arXiv:1610.07390v1 [cs.PL] 24 Oct 2016

A Practical Approach to Interval Refinement for `math.h/cmath` Functions

Roberto Bagnara[1,2], Michele Chiari[2,3], Roberta Gori[4], and Abramo Bagnara[2]

[1] Dipartimento di Matematica e Informatica, University of Parma, Italy
bagnara@cs.unipr.it
[2] BUGSENG srl, http://bugseng.com, Italy
roberto.bagnara@bugseng.com
[3] University of Parma, Italy
michele.chiari1@studenti.unipr.it
[4] Dipartimento di Informatica, University of Pisa, Italy
gori@di.unipi.it

Abstract. Verification of C/C++ programs has seen considerable progress in several areas, but not for programs that use these languages' mathematical libraries. The reason is that all libraries in widespread use come without no guarantees about the computed results. This would seem to prevent any attempt at formal verification of programs that use them: without a specification for the functions, no conclusion can be statically drawn about the behavior of the program. We propose an alternative to surrender: even if we do not have a proper specification for the functions, we might have a partial specification; at the very least we have their implementation, for which a partial specification can sometimes be extracted. When even a partial specification is unavailable, we can still detect flaws in the program and facilitate its partial verification with testing. We propose a pragmatic, practical approach that leverages the fact that most `math.h/cmath` functions are *almost* piecewise monotonic: they may have *glitches*, often of very small size and in small quantities. We develop interval refinement techniques for such functions that enable verification via abstract interpretation, symbolic model checking and constraint-based test data generation.

1 Introduction

The use of floating-point computations for the implementation of critical systems is perceived as increasingly acceptable. Even in modern avionics, one of the most critical domains for software, floating-point numbers are now used, more often than not, instead of fixed-point arithmetic [BDL12].

Acceptance of floating-point computations in the design of critical systems was facilitated by the wide adoption of significant portions of the IEEE 754 standard for binary floating-point arithmetic [IEE08]. Most implementations of the C and C++ programming languages provide floating-point data types that conform to IEEE 754 as far as the basic arithmetic operations and the conversions are concerned. Some implementations push compliance further by providing other

operations required by IEEE 754, such as correctly rounded *fused multiply-add*[5] and square root [IEE08, Section 5].

The C and C++ floating-point mathematical functions are part of the standard libraries of the respective languages. Different versions of the language standards define different, though backwards-compatible, sets of basic mathematical functions (see, e.g., [Int11,Int14]). Access to such functions requires inclusion of the `math.h` header file, in C, or the `cmath` header file in C++. While the IEEE 754 provides *recommendations* for these functions [IEE08, Section 9.2], no C/C++ implementation we know of complies to the recommendation that such functions be correctly rounded. An exception is CR-LIBM,[6] a truly remarkable research prototype that, however, has still not found industrial adoption, possibly because of the worst-case performance of some functions (the average case being usually very good). Another freely available library of correctly-rounded functions is libmcr,[7] by Sun Microsystems, Inc. It provides some double-precision transcendental elementary functions (`exp`, `log`, `pow`, `sin`, `cos`, `tan`, `atan`), but its development stopped in 2004. Even though we cannot exclude the existence of proprietary implementations of `libm` providing formalized precision guarantees, we were not able to find one.

For comparison, Numeric Annex (G) of the Ada Reference Manual [TDB+06] prescribes that, when the result of the evaluation of an elementary function does not result into an exception, the numerical result belongs to an interval defined as the smallest floating-point interval in the format that contains all the values of the form $f \cdot (1 + d)$, where

- f is the exact value of the corresponding mathematical function at the given parameter values;
- d is a real number such that $|d|$ is less than or equal to the function's maximum relative error.

Each elementary function is then given with a maximum relative error. This is expressed in terms of the "relative error of the format", E_{rel}, defined as the smallest positive number in the format satisfying $1.0 \oplus E_{\text{rel}} \neq 1.0$, where $\oplus$ denotes floating-point addition. Then:

- `sqrt`, `sin` and `cos` have a maximum relative error of $2 \cdot E_{\text{rel}}$ (but see below for the restrictions concerning periodic functions);
- `log`, `exp`, `tan` and `cot` have a maximum relative error of $4 \cdot E_{\text{rel}}$;
- the forward and inverse hyperbolic functions have a maximum relative error of $8 \cdot E_{\text{rel}}$;
- the maximum relative error for exponentiation `x**y` depends on the parameter values x and y: $\big(4 + |y \log(x)|/32\big) \cdot E_{\text{rel}}$.

[5] Fused multiply-add applied to x, y and z yields $x \cdot y + z$ rounded as per the current rounding mode.

[6] See `http://lipforge.ens-lyon.fr/www/crlibm/`, last accessed on October 15th, 2016.

[7] See `https://github.com/simonbyrne/libmcr/blob/master/README`, last accessed on October 15th, 2016.

Note, though, that the maximum relative error given above for the forward trigonometric functions applies only when the absolute value of the angle parameter is less than or equal to some implementation-defined angle threshold. Finally, the following specifications also take precedence over the maximum relative error bounds:

- the absolute value of the result of the `sin`, `cos`, and `tanh` functions never exceeds one;
- the absolute value of the result of the `coth` functions is never less than one;
- the result of the `cosh` functions is never less than one.

Summarizing, differently from the C and the C++ standard, the Ada standard gives some guarantees about the result of mathematical functions, but such a partial specification is still problematic from the point of view of formal verification. In the C/C++ world, the only implementation we know that says something about errors is GNU libc [LSM+16]:

> *"Therefore many of the functions in the math library have errors. The table lists the maximum error for each function which is exposed by one of the existing tests in the test suite. The table tries to cover as much as possible and list the actual maximum error (or at least a ballpark figure) but this is often not achieved due to the large search space."*

So, nothing that can really be trusted in a safety-critical context. In the embedded world, we checked the documentation of all major toolchain providers and for two of them we found that the lack of guarantees is explicit mentioned, whereas for all the others nothing is said in this regard. It is to mitigate this state of affairs that test-suites like *C11/C99 FPCE Test Suite* by Tydeman Consulting[8] have been created. Solid Sands' *SuperTest*,[9] one of the leading test and validation suites for C and C++ compilers, has been recently extended with many more tests for the mathematical libraries.

To illustrate the concrete problem raised by the use of transcendental functions in program verification settings, consider the code reproduced in Listing 1.1. It is a reduced version of a real-world example extracted from a critical embedded system.[10] Some of the questions to be answered for this code are:

1. Can infinities and NaNs be generated?
2. Can `sin`, `cos`, and `tan` be invoked on ill-conditioned arguments?
3. If so, which inputs to the given functions may cause such anomalies?

Concerning the second question, we call the argument of a floating-point trigonometric function *ill-conditioned* if its absolute value exceeds some application-dependent threshold. Ideally, this threshold should be just above π. To understand this often-overlooked programming error, consider that, if x is an IEEE 754

[8] `http://www.tybor.com/`, last accessed on October 23rd, 2016.

[9] `http://www.solidsands.nl/`, last accessed on October 23rd, 2016.

[10] The original source code is available at `http://paparazzi.enac.fr`, Paparazzi UAV (Unmanned Aerial Vehicle), v5.10.0_stable release, file `sw/misc/satcom/tcp2ivy.c`, last accessed on October 15th, 2016.

Listing 1.1. Code excerpted from a real-world avionic library

```
#include <math.h>
#include <stdint.h>

#define RadOfDeg(x) ((x) * (M_PI/180.))
#define E 0.08181919106 /* Computation for the WGS84 geoid only */
#define LambdaOfUtmZone(utm_zone) RadOfDeg((utm_zone-1)*6-180+3)
#define CScal(k, z) { z.re *= k; z.im *= k; }
#define CAdd(z1, z2) { z2.re += z1.re; z2.im += z1.im; }
#define CSub(z1, z2) { z2.re -= z1.re; z2.im -= z1.im; }
#define CI(z) { float tmp = z.re; z.re = - z.im; z.im = tmp; }
#define CExp(z) { float e = exp(z.re); z.re = e*cos(z.im); \
  z.im = e*sin(z.im); }
#define CSin(z) { CI(z); struct complex _z = {-z.re, -z.im}; \
  float e = exp(z.re); float cos_z_im = cos(z.im); z.re = e*cos_z_im; \
  float sin_z_im = sin(z.im); z.im = e*sin_z_im; _z.re = cos_z_im/e; \
  _z.im = -sin_z_im/e; CSub(_z, z); CScal(-0.5, z); CI(z); }

static inline float isometric_latitude(float phi, float e) {
  return log(tan(M_PI_4 + phi / 2.0))
    - e / 2.0 * log((1.0 + e * sin(phi)) / (1.0 - e * sin(phi)));
}

static inline float isometric_latitude0(float phi) {
  return log(tan(M_PI_4 + phi / 2.0));
}

void latlong_utm_of(float phi, float lambda, uint8_t utm_zone) {
  float lambda_c = LambdaOfUtmZone(utm_zone);
  float ll = isometric_latitude(phi, E);
  float dl = lambda - lambda_c;
  float phi_ = asin(sin(dl) / cosh(ll));
  float ll_ = isometric_latitude0(phi_);
  float lambda_ = atan(sinh(ll) / cos(dl));
  struct complex z_ = { lambda_, ll_ };
  CScal(serie_coeff_proj_mercator[0], z_);
  uint8_t k;
  for(k = 1; k < 3; k++) {
    struct complex z = { lambda_, ll_ };
    CScal(2*k, z);
    CSin(z);
    CScal(serie_coeff_proj_mercator[k], z);
    CAdd(z, z_);
  }
  CScal(N, z_);
  latlong_utm_x = XS + z_.im;
  latlong_utm_y = z_.re;
}
```

single-precision number and $x \geq 2^{23}$, then the smallest single-precision range containing $[x, x + 2\pi)$ contains no more than three floating-point numbers.[11] Current implementations of floating-point trigonometric functions contain precise range reduction algorithms that have no problem computing a very accurate result even for numbers much higher than 2^{23}: the point is that the input is so inaccurate that the very accurate result is, from the point of view of the application, indistinguishable from a pseudo-random number.

We do not have a precise specification for the library functions that are assumed in the code of Listing 1.1. While the project it comes from is said to have been designed from the beginning with portability in mind,[12] the sources and build machinery refer to GNU libc and to Newlib,[13] but no specific versions are mentioned. We can probably assume assume a POSIX-compliant behavior with respect to special values, e.g.: if `log()` is called with a negative number, a NaN is returned; if `atan()` is called with ± 1 then an infinity is returned [IEE13]. This information is not sufficient in order to provide a general answer to the verification questions asked above. Things change if, in addition to assuming a compilation mode that strictly adheres to IEEE 754, we fix a specific implementation of the mathematical library.

In this paper, we propose a pragmatic, practical approach that enables verification of C/C++ programs using `math.h`/`cmath` functions, even with minimal or no specification in addition to the special cases mandated by the standards the implementation conforms to, e.g., POSIX [IEE13]. We observed that, for all the implementations of `libm` we tested, the piecewise monotonicity property of the corresponding real functions is almost preserved. Consider, for instance the `tanhf()` function, which is meant to approximate the hyperbolic tangent function over IEEE 754 single-precision floats. While $y = \tanh(x)$ is monotonically (strictly) increasing over $(-\infty, +\infty)$, `y = tanhf(x)` can be monotonically non-decreasing over the full range of IEEE 754 single-precision floats, for some rounding mode, or it can be "almost monotonically non-decreasing." By this we mean that, going from $-\infty$ to $+\infty$, there may be an occasional drop in the graph of `tanhf()`, but this is quickly recovered from, that is, the function starts increasing again. We use the term *glitches* to name such occasional drops: we observed that glitches are often shallow (most often just one ULP), narrow (most often just 2 ULPs), and, on average, not very frequent. We leverage this fact to provide general interval refinement algorithms that enable software verification with various techniques in all cases:

– If we have approximate but correct information about the maximal depth and width of glitches and, possibly, their number and their localization, then we can guarantee containment for the refined intervals, and this allows us to perform formal verification via abstract interpretation, symbolic model checking or automatic theorem proving:

[11] Substitute 2^{23} with 2^{52} and the same holds for IEEE 754 double-precision numbers.
[12] See `http://wiki.paparazziuav.org/`, last accessed on October 20th, 2016.
[13] See `https://sourceware.org/newlib/`, last accessed on October 20th, 2016.

  - with precise information and small/few glitches (which includes the case where there are no glitches at all, which includes the case of correctly rounded implementations), the refinement will result in tight intervals, verification will be computationally cheaper and with less "don't knows";
  - with less precise information, verification is still possible, at a higher cost and possibly more "don't knows".
- With possibly incorrect information about the presence and nature of glitches, we can still automatically generate test input data. Things can go wrong in two ways:
  - the refined intervals are too tight, in which case we might fail to generate test inputs for paths and conditions that are indeed feasible;
  - the refined intervals are too lax, in which case we might generate test inputs that fail to exercise the paths or conditions they are supposed to exercise.

  In both cases, this is not a big deal if we are still able to generate lots of useful and accurate tests.

The plan of the paper is as follows: Section 2 recalls basic definitions and introduces the required notation; Section 3 introduces the notions of *monotonicity glitch* and *quasi-monotonicity*; Section 4 explains how trigonometic functions, which are periodic, can be treated by partitioning a subset of their graph into a set of quasi-monotonic branches; Section 5 describes direct and indirect propagation algorithms that are able to deal with (at least) 75 of the `math.h`/`cmath` functions; Section 6 briefly describes the implementation in the context of the ECLAIR software verification platforms and illustrates the result of an experiment; Section 7 discusses the problems that remain to be solved and sketches several ideas for future work; Section 8 concludes the main part of the paper. Appendix A presents additional data about the glitches in the single-precision functions for several implementations of the `math.h`/`cmath` mathematical functions; Appendix B presents formal proofs of correctness of the most complex algorithms of Section 5.

2 Preliminaries

We will denote by $\mathbb{R}_+$ and $\mathbb{R}_-$ the sets of strictly positive and strictly negative real numbers, respectively.

Definition 1. (IEEE 754 binary floating-point numbers.) *A set of IEEE 754 binary floating-point numbers* [IEE08] *is uniquely identified by:* $p \in \mathbb{N}$*, the number of significant digits (precision);* $e_{\max} \in \mathbb{N}$*, the maximum exponent;* $e_{\min} \in \mathbb{N}$*, the minimum exponent (this can be either* $1 - e_{\max}$ *or* $-e_{\max}$*, the former being preferred). The set of binary floating-point numbers* $\mathbb{F}(p, e_{\max}, e_{\min})$ *includes:*

- *all signed zero and non-zero numbers of the form* $(-1)^s \cdot 2^e \cdot m$*, where*
  - s *is the* sign bit*;*
  - *the* exponent e *is any integer such that* $e_{\min} \leq e \leq e_{\max}$*;*

- *the* mantissa m*, with* $0 \leq m < 2$*, is a number represented by a string of* p *binary digits with a "binary point" after the first digit:*

$$m = (d_0 \,.\, d_1 d_2 \ldots d_{p-1})_2 = \sum_{i=0}^{p-1} d_i 2^{-i};$$

– *the* infinities $+\infty$ *and* $-\infty$*.*

The smallest positive normal *floating-point number is* $f_{\min}^{\text{nor}} := 2^{e_{\min}}$ *and the largest is* $f_{\max} := 2^{e_{\max}}(2 - 2^{1-p})$*. The non-zero floating-point numbers whose absolute value is less than* $2^{e_{\min}}$ *are called* subnormal*: they always have fewer than* p *significant digits. Every finite floating-point number is an integral multiple of the smallest subnormal magnitude* $f_{\min} := 2^{e_{\min}+1-p}$*. Note that the* signed zeroes $+0$ *and* -0 *are distinct floating-point numbers.*

In the sequel we will only be concerned with IEEE 754 binary floating-point numbers and we will write simply $\mathbb{F}$ for $\mathbb{F}(p, e_{\max}, e_{\min})$ when there is no risk of confusion.

Definition 2. (Floating-point symbolic order.) *Let* $\mathbb{F}$ *be any IEEE 754 floating-point format. The relation* $\prec \subseteq \mathbb{F} \times \mathbb{F}$ *is such that, for each* $x, y \in \mathbb{F}$*,* $x \prec y$ *if and only if either:* $x = -\infty$ *and* $y \neq -\infty$*, or* $x \neq +\infty$ *and* $y = +\infty$*, or* $x = -0$ *and* $y \in \{+0\} \cup \mathbb{R}_+$*, or* $x \in \mathbb{R}_- \cup \{-0\}$ *and* $y = +0$*, or* $x, y \in \mathbb{R}$ *and* $x < y$*. The function* $(\cdot)^+ \colon \big(\mathbb{F} \setminus \{+\infty\}\big) \to \mathbb{F}$ *is defined, for each* $x \in \mathbb{F} \setminus \{+\infty\}$*, by* $x^+ := \min\{\, y \in \mathbb{F} \mid x \prec y \,\}$*. Similarly,* $(\cdot)^- \colon \big(\mathbb{F} \setminus \{-\infty\}\big) \to \mathbb{F}$ *is defined, for each* $y \in \mathbb{F} \setminus \{-\infty\}$*, by* $y^- := \max\{\, x \in \mathbb{F} \mid x \prec y \,\}$*. We will iteratively apply these functions, so that, e.g., for each* $n \in \mathbb{N}$*, we will refer to the partial function* $(\cdot)^{+^n} \colon \mathbb{F} \rightarrowtail \mathbb{F}$ *given, for each* $x \in \mathbb{F}$*, by*

$$\begin{cases} x^{+^0} := x; \\ x^{+^{n+1}} := \big(x^{+^n}\big)^+, \text{ if } x^{+^n} \neq +\infty; \end{cases}$$

The partial order $\preccurlyeq \subseteq \mathbb{F} \times \mathbb{F}$ *is such that, for each* $x, y \in \mathbb{F}$*,* $x \preccurlyeq y$ *if and only if either* $x \prec y$ *or* $x = y$*. We use the notation* $x \succ^n y$ *to signify that* $x \succ y^{+^n}$*.*

Note that $\mathbb{F}$ is linearly ordered with respect to '$\prec$'.

For $x \in \mathbb{F}$, we will sometimes confuse the floating-point number with the extended real number it represents, the floats -0 and $+0$ both corresponding to the real number 0. Thus, when we write, e.g., $x < y$ we mean that x is numerically less than y (for example, we have $-0 \prec +0$ although $-0 \not\prec +0$, but note that $x \preccurlyeq y$ implies $x \leq y$ if x and y are finite).

Definition 3. (Floating-point intervals.) *Let* $\mathbb{F}$ *be any IEEE 754 floating-point format. The set* $\mathcal{I}_{\mathbb{F}}$ *of floating-point intervals with boundaries in* $\mathbb{F}$ *is*

$$\mathcal{I}_{\mathbb{F}} := \{\varnothing\} \cup \big\{\, [l, u] \;\big|\; l, u \in \mathbb{F}, l \preccurlyeq u \,\big\}.$$

$[l, u]$ *denotes the set* $\{\, x \in \mathbb{F} \mid l \preccurlyeq x \preccurlyeq u \,\}$*.* $\mathcal{I}_{\mathbb{F}}$ *is a bounded meet-semilattice with least element* $\varnothing$*, greatest element* $[-\infty, +\infty]$*, and the meet operation, which is induced by set-intersection, will be simply denoted by* $\cap$*.*

Floating-point intervals with boundaries in $\mathbb{F}$ allow to capture the extended numbers in $\mathbb{F}$: NaNs should be tracked separately.

Given an interval $[l, u] \in \mathcal{I}_{\mathbb{F}}$ we denote by $\#[l, u]$ the cardinality of the set $\{\, x \in \mathbb{F} \mid l \preccurlyeq x \preccurlyeq u \,\}$.

3 (Quasi-) Monotonicity and Glitches

A monotonic real-valued partial function $\hat{f} \colon \mathbb{R} \rightarrowtail \mathbb{R}$ is called *monotonic* if it is order preserving (i.e., $\hat{f}(x) \leq \hat{f}(y)$ whenever $x \leq y$ and both $\hat{f}(x)$ and $\hat{f}(y)$ are defined), in which case we call it *isotonic*, or if it is order reversing (i.e., $\hat{f}(x) \geq \hat{f}(y)$ whenever $x \leq y$ and both $\hat{f}(x)$ and $\hat{f}(y)$ are defined), in which case we say $\hat{f}$ is *antitonic*.

Definition 4. (Quasi-monotonicity.) *Let $f \colon \mathbb{F} \to \mathbb{F}$ be a floating-point function meant to approximate a real-valued partial function $\hat{f} \colon \mathbb{R} \rightarrowtail \mathbb{R}$. Let also $I \subseteq \mathbb{F}$ be a floating-point interval. We say that f* is quasi-monotonic / isotonic / antitonic on I *if $\hat{f}$ is always defined and monotonic / isotonic / antitonic on I.*

Let $f \colon \mathbb{F} \to \mathbb{F}$ be a quasi-monotonic function. The best we can hope for is that f be monotonic over $(\mathbb{F}, \preccurlyeq)$ for all rounding modes. While this is often the case, it is not always the case: monotonicity is occasionally violated at spots we call *monotonicity glitches*.

Definition 5. (Monotonicity glitches.) *Let $f \colon \mathbb{F} \to \mathbb{F}$ be a quasi-isotonic function on $I \subseteq \mathbb{F}$. An* isotonicity glitch of f in I *is an interval $[l, u] \subseteq I$ such that:*

$$u \succ l^{+} \quad \wedge \quad \forall x \in (l, u) : f(l) \succ f(x) \quad \wedge \quad f(l) \preccurlyeq f(u).$$

If f is quasi-antitonic, an antitonicity glitch of f in I *is an isotonicity glitch of $-f$ in I. Isotonicity and antitonicity glitches are collectively called* monotonicity glitches *or, simply,* glitches.

Let $G = [l, u]$ be a monotonicity glitch of f in I. The width *and the* depth *of G are given, respectively, by*

$$\begin{aligned} &\operatorname{width}(G) := \#[l, u] - 1, \\ &\operatorname{depth}(G) := \#\big[m, f(u)\big] - 1, \qquad\qquad \text{where } m = \min_{x \in (l,u)} f(x). \end{aligned}$$

Note that, for each glitch G, we have $\operatorname{width}(G) \geq 2$ *and* $\operatorname{depth}(G) \geq 1$. *A glitch of f in I is called* maximal *if none of its proper supersets is a glitch of f in I. Non-maximal glitches are also called* sub-glitches.

See Figure 1 for an exemplification of these concepts: G_1 is a maximal glitch; G_2, being contained into G_1 is non-maximal; $\operatorname{width}(G_1) = 5$, $\operatorname{depth}(G_1) = 4$, $\operatorname{width}(G_2) = \operatorname{depth}(G_2) = 2$;

Tables 1–3 show the relevant statistics about glitches for 25 functions provided by `libm` on several implementations. The tested implementations are introduced in Table 1: for each implementation the table lists its identification

Table 1. Glitch data: main characteristics of the tested implementations

id	CPU	OS	compiler	`libm` version
`alpha`	x86_64	Ubuntu 14.04	GCC 4.8.4	EGLIBC 2.19
`gcc110`	POWER7	Fedora 20	GCC 4.8.1	GNU libc 2.18
`gcc111`	POWER7	AIX 7	GCC 4.8.1	
`gcc112`	POWER7	Fedora 21	GCC 4.9.2	GNU libc 2.20
`gcc113`	AArch64	Ubuntu 14.04	GCC 4.8.4	EGLIBC 2.19
`igor`	x86_64	Fedora 12	GCC 4.4.4	GNU libc 2.11.2
`raspi`	ARMv6 + VFPv2	Raspbian Jessie	GCC 4.9.2	GNU libc 2.19
`zoltan`	x86_64	Ubuntu 16.04	GCC 5.4.0	GLIBC 2.23

Table 2. Glitch data for the `gcc110/2/3` machines

function	$D_{\min}$	D_{M}	near			up			down			zero		
			n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}
`acosf`	-1	1												
`acoshf`	1	∞							1	1	2	1	1	2
`asinf`	-1	1												
`asinhf`	$-\infty$	∞							2	1	2	2	1	2
`atanf`	$-\infty$	∞				1	1	10^8						
`atanhf`	-1	1							2	1	2	2	1	2
`cbrtf`	$-\infty$	∞	10^6	1	2	10^6	1	2	10^6	1	2	10^6	1	2
`coshf`	$-\infty$	∞	456	1	2	462	1	2	442	1	2	448	1	2
`erff`	$-\infty$	∞												
`expf`	$-\infty$	∞												
`exp10f`	$-\infty$	∞												
`exp2f`	$-\infty$	∞	2	1	2									
`expm1f`	$-\infty$	∞												
`lgammaf`	2	∞												
`logf`	0	∞												
`log10f`	0	∞												
`log1pf`	-1	∞							1	1	2	1	1	2
`log2f`	0	∞												
`sinhf`	$-\infty$	∞												
`sqrtf`	0	∞												
`tanhf`	$-\infty$	∞				1	1	2	2	1	3			
`tgammaf`	2	∞	10^4	2	3	10^4	4	4	10^4	2	3	10^4	3	4
`cosf`	-2^{23}	2^{23}	10^4	1	3	10^4	1	3	10^4	1	3	10^4	1	3
`sinf`	-2^{23}	2^{23}												
`tanf`	-2^{23}	2^{23}												

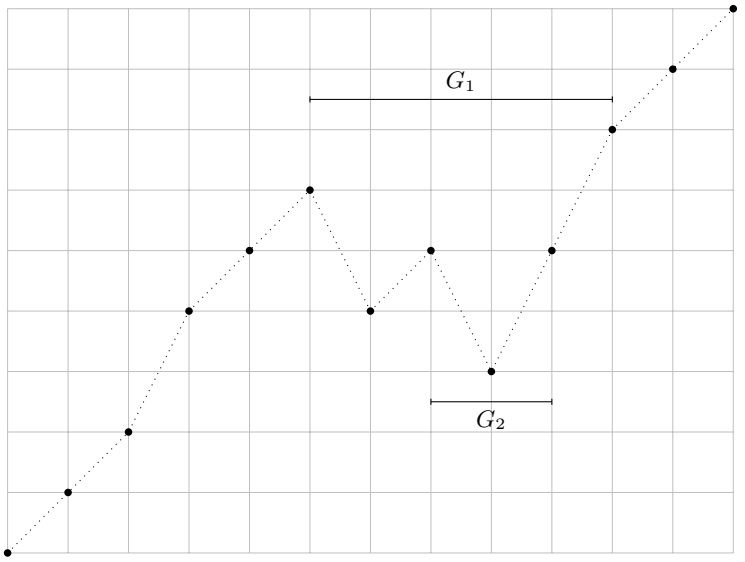

Fig. 1. An example of monotonicity glitches

code, which is used in the other tables, the CPU architecture, the operating system, and, where known, the `libm` version.

The data collected for two implementations are presented in Tables 3–3; data for the other implementations are reported in Appendix A. These tables present, for each function, its name and the minimum and maximum of the considered domain interval. For most of the functions, such interval is the natural one. The exceptions are the following: for `lgammaf` we start at 2, which is where monotonicity theoretically begins; for `tgammaf` we also start at 2 because the considered implementations behave in a bizarre way for arguments less than 2; for the trigonometric functions, at the bottom of the tables, we restrict the domain to a region where there are at least 12 floats per period, for reasons that will be discussed in Section 4.1.

For each function and each rounding mode (near, up, down, zero), Tables 4–6 give the number of glitches, n_{g}, their maximum depth, d_{M}, and their maximum width, w_{M}. For the trigonometric functions we report the cumulative results concerning all the quasi-isotonic and quasi-antitonic branches in the given range. In the columns labeled n_{g}, we report for these functions the maximum number of glitches in any such quasi-monotonic branch (we will refer to this quantity in Section 5.3 as n_{gM}).

The following observations can be made:

1. there are few glitches: many functions have no glitch at all, several functions have just a few glitches, a few functions have many glitches;
2. most glitches are very shallow;

Table 3. Glitch data for the `zoltan` machine

function	$D_{\min}$	D_{M}	near			up			down			zero		
			n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}
`acosf`	-1	1												
`acoshf`	1	∞							1	1	2	1	1	2
`asinf`	-1	1												
`asinhf`	$-\infty$	∞							2	1	2	2	1	2
`atanf`	$-\infty$	∞				1	1	10^8						
`atanhf`	-1	1							2	1	2	2	1	2
`cbrtf`	$-\infty$	∞	10^6	1	2	10^6	1	2	10^6	1	2	10^6	1	2
`coshf`	$-\infty$	∞	454	1	2	466	1	2	442	1	2	448	1	2
`erff`	$-\infty$	∞												
`expf`	$-\infty$	∞												
`exp10f`	$-\infty$	∞												
`exp2f`	$-\infty$	∞	1	1	2	3	1	2	2	1	2	1	1	2
`expm1f`	$-\infty$	∞												
`lgammaf`	2	∞	163	1	2	164	1	2	166	1	2	161	1	2
`logf`	0	∞												
`log10f`	0	∞												
`log1pf`	-1	∞							1	1	2	1	1	2
`log2f`	0	∞												
`sinhf`	$-\infty$	∞												
`sqrtf`	0	∞												
`tanhf`	$-\infty$	∞				1	1	2	2	1	3			
`tgammaf`	2	∞	10^5	4	3	10^5	4	3	10^5	4	3	10^5	4	3
`cosf`	-2^{23}	2^{23}												
`sinf`	-2^{23}	2^{23}												
`tanf`	-2^{23}	2^{23}												

3. with a notable exception, glitches are also very narrow.

It is important to observe that glitches are not simply bugs that will surely be fixed at the next release. For instance, the implementation of `tgammaf` of Ubuntu 16.04/x86_64 has more numerous and deeper glitches than the one in Ubuntu 14.04/x86_64. The point is that monotonicity is not one of the objective of most implementations of `math.h`/`cmath` functions. For instance, both the manual [LSM+16] and the FAQ of GNU libc explicitly exclude monotonicity from the accuracy goals of the library, so that bug reports about violated monotonicity are closed as invalid.[14]

We will now deal with the special case of the trigonometric functions, as they pose the additional problem of periodic slope inversions. Then we will see how quasi-monotonicity can be exploited for the purposes of interval refinement and, in turn, of software verification.

[14] See, e.g., bug reports `https://sourceware.org/bugzilla/show_bug.cgi?id=15898` and `https://sourceware.org/bugzilla/show_bug.cgi?id=15899`.

4 Trigonometric Functions

Treating glitches in periodic trigonometric functions (`sinf`, `cosf` and `tanf`) requires a somewhat more complex approach. Since the underlying theoretical (partial) functions in $\mathbb{R} \rightarrowtail \mathbb{R}$ change their monotonicity periodically, Definition 5 would fail to distinguish between actual glitches caused by the implementation and the said "legitimate" monotonicity changes.

Moreover, we must consider the fact that the distance between two consecutive floating point numbers becomes wider as the exponent becomes bigger, while the period of trigonometric functions remains the same.

In order to properly identify and characterize monotonicity glitches in the trigonometric functions, we need to find

- an interval in which the amount of floating point numbers in every quasi-monotonic branch of the functions' graphs is sufficiently high;
- a way to distinguish glitches from legitimate changes in function monotonicity in the said interval.

4.1 An Appropriate Domain for Trigonometric Functions Analysis

When choosing a suitable domain for the search of monotonicity glitches, we must carefully take into account the value of the ulp: $\mathbb{R} \to \mathbb{R}$ function in that interval: we will use the definition of ulp given in [Mul05, Definition 5]. A floating-point interval in which the idea of monotonicity glitch is well defined should have a sufficient cardinality to allow for glitches that do not cover the entire interval.

Let $x \in \mathbb{F}$ be a positive normal floating-point number: then $\mathrm{ulp}(x) = x - x^- = 2^{e_{x^-} - p + 1}$, where p is the precision of the format, and e_{x^-} the exponent of x^-. For the considered trigonometric functions, the size in $\mathbb{R}$ of those intervals in which the theoretical function is constantly monotonic or antitonic is π. If we consider an interval $[-\ell_{\max}, \ell_{\max}]$ in which for each $x \in [-\ell_{\max}, \ell_{\max}]$ we have $\mathrm{ulp}(x) \leq 0.5$, then each monotonic branch contains at least 6 or 7 floats, which is acceptable to be used with the propagation algorithms described in Section 5. In intervals with a higher ulp value, the definition of monotonicity glitches would be hardly meaningful.

Therefore, a maximum exponent $e_{\ell_{\max}} = p - 1$ will be used, leading to a domain $[-\ell_{\max}, \ell_{\max}]$ with $\ell_{\max} = 2^{p-1}$. In conclusion, for the IEEE 754 Single Precision format, the value of $\ell_{\max}$ will be 2^{23}, and for the Double Precision format $\ell_{\max} = 2^{52}$.

4.2 Glitches in Trigonometric Functions

With the already stated definitions of glitches, the normal changes in monotonicity of trigonometric functions would be detected as glitches. Actual glitches would be recognized as sub-glitches of these intervals, thus being difficult to discern.

On the other hand, Definition 5 is still useful if we separately consider those intervals in which the functions are quasi-isotonic or quasi-antitonic. The following definitions let us denote appropriate intervals in which the functions have such properties. Definition 6 describes the maximal sets of intervals in which the function has a constant monotonicity, while Definition 7 applies the definition of glitches to those intervals.

Definition 6. (Set of the quasi-isotonic (quasi-antitonic) intervals.) *Let $f\colon \mathbb{R} \to \mathbb{R}$ be a periodic function, and let $\mathcal{M}_f^{\mathbb{R}}$ ($\mathcal{A}_f^{\mathbb{R}}$) be the set of the maximal isotonic (resp. antitonic) intervals of f in $\mathbb{R}$:*

$$\mathcal{M}_f^{\mathbb{R}} = \{I \in \mathcal{I}_{\mathbb{R}} \mid f \textit{ is } \text{isotonic } \textit{on } I \\ \wedge \forall I' \in \mathcal{I}_{\mathbb{R}} : I \subset I' \implies f \textit{ is } \text{not isotonic } \textit{on } I'\}.$$

Let $D \subseteq \mathbb{F}$ such that, for each $x \in D$, $f(x) \in \mathbb{R}$. The set of the quasi-isotonic intervals of f in D is a set $\mathcal{M}_f^D \subseteq \mathcal{I}_{\mathbb{F}}$ (Definition 3) such that:

$$\mathcal{M}_f^D = \{F \in \mathcal{I}_{\mathbb{F}} \mid F \subseteq D \wedge \exists I \in \mathcal{M}_f^{\mathbb{R}} \,.\, F = I \cap D\}$$

When an interval $I \in \mathcal{M}_f^{\mathbb{R}}$ has 0 as the lower or upper bound, the corresponding interval in $\mathcal{M}_f^D$ has $+0$ or -0 as the corresponding bound.

The sets of the (maximal) quasi-antitonic intervals $\mathcal{A}_f^{\mathbb{R}}$ and $\mathcal{A}_f^D$ are defined similarly.

Definition 7. (Isotonicity glitch in a periodic function.) *Let $f\colon \mathbb{F} \to \mathbb{F}$ be the implementation of a periodic function, let $I \subseteq \mathbb{F}$ such that, for each $x \in I$, $f(x) \in \mathbb{R}$, and let $\mathcal{M}_f^I$ be the set of the isotonic intervals of f in I. An* isotonicity glitch of f in I *is an interval $G \subseteq I$ such that:*

1. *there exists an interval $I_{\mathcal{M}} \in \mathcal{M}_f^I$ such that $G \subseteq I_{\mathcal{M}}$;*
2. *G is an isotonicity glitch of f in $I_{\mathcal{M}}$ according to Definition 5.*

Antitonicity glitches in periodic functions are defined similarly, but substituting $\mathcal{M}_f^I$ with $\mathcal{A}_f^I$. Glitch *width* and *depth* in periodic functions are consequently defined according to Definition 5.

4.3 Range Reduction for Trigonometric Functions

In order to split the domain of trigonometric functions into quasi-monotonic intervals, algorithms capable of discerning close floating-point approximations of multiples of $\frac{\pi}{2}$ are needed. Such algorithms will be very similar to the range reduction algorithms used to implement trigonometric library functions [PH83] [DMMM96,Ng92,BKD+05]. Let $x \in \mathbb{F}$ be the argument of a trigonometric function: argument reduction is aimed at finding values $k \in \mathbb{Z}$ and $r \in \mathbb{R}$ such that $x = k\frac{\pi}{2} + r$. In order to obtain these values, $k = \left\lfloor x\frac{2}{\pi} \right\rfloor$ and an approximation of $r = \frac{\pi}{2}\left(x\frac{2}{\pi} - k\right)$ are computed with an appropriate increased precision. Then, r is used to compute the value of the function in a reduced interval, and the

last digits of k are used to identify the quadrant of x. A notable implementation of range reduction algorithms is the one in `fdlibm`[15]; it is also inluded in some other `libm` implementations, such as GNU libc[16], which uses it for some architectures. Other important implementations are the ones of the already mentioned CR-LIBM and libmcr.

However, as we shall point out, the algorithms needed for our analysis do not have the same precision requirements as actual range reduction algorithms: since we will not actually compute the functions, r is not needed in full precision. An algorithm capable of computing integer division by $\frac{\pi}{2}$ is needed, as well as a generator of floating-point approximations of $\frac{\pi}{2}$.

The computations of such operations poses the problem of deciding how many digits of the approximation of $\frac{\pi}{2}$ are needed in order to obtain correctly rounded results. This is also known as the *Table Maker's Dilemma* [LMT98]. Lemma 1 allows us to answer this question, by computing the rounding error generated by the multiplication of a floating point number k, which shall be considered exact, and a floating-point approximation $[x]$ of a real number, x. Any of the rounding modes defined in the IEEE 754 standard can be used to generate the approximation, and to perform the multiplication.

Lemma 1. *Let $k \in \mathbb{F}$ be a number in any of the IEEE 754 binary floating-point formats, with an exponent e_k. Let $x \in \mathbb{R}$, such that $\nexists n \in \mathbb{Z} \ . \ |x| = 2^n$ and $f_{\min} < |x| < f_{\max}$. Let $[x] \in \mathbb{F}$ be an approximation of x with p significant digits and exponent e_x.*

Then, the absolute error of the correctly rounded floating-point multiplication between k and $[x]$ is $|\nabla(k \boxdot [x])| \leq 2^{e_k + e_x - p + 4}$.

Proof. See Section B.3 of Appendix B.

The algorithms for the propagation of constraints with trigonometric functions, which will be presented in Section 5.3, divide the domains of trigonometric functions into quasi-monotonic intervals. They need functions to identify approximations of multiples of $\frac{\pi}{2}$ (by computing $\left\lceil x\frac{2}{\pi}\right\rceil$, $x \in \mathbb{F}$), and to generate such upper and lower approximations ($\left[k\frac{\pi}{2}\right]_{\uparrow}$ and $\left[k\frac{\pi}{2}\right]_{\downarrow}$, $k \in \mathbb{Z}$) in the floating-point format in which the studied trigonometric functions operate, which will be called the "target" format.

Integer division by $\frac{\pi}{2}$, rounded upwards This procedure has the purpose of computing a number k such that $(k-1)\frac{\pi}{2} \leq x \leq k\frac{\pi}{2}$. It does so by multiplying x by $\frac{2}{\pi}$, and then rounding the result to an integer towards positive infinity. The major caveat in this procedure is the precision at which the multiplication should be performed, and consequently the number of digits of $\frac{2}{\pi}$ that should be stored. For $x \times \frac{2}{\pi}$ to be rounded to the correct integer, we need a precision p

[15] Developed in 1993 by Sun Microsystems, Inc. See `http://www.netlib.org/fdlibm/`, file `e_rem_pio2.c`. Last accessed on October 17th, 2016.

[16] See GNU libc 2.22, file `sysdeps/ieee754/dbl-64/e_rem_pio2.c`.

sufficient to preserve at least the sign and order of magnitude of the difference $\Delta x = x\frac{2}{\pi} - \hat{k}$, where $\hat{k}$ is the integer nearest to $x\frac{2}{\pi}$:

$$\left| \nabla \left(x \boxdot \left[\frac{2}{\pi} \right] \right) \right| < 2^{e_{\Delta x}},$$

where $e_{\Delta x}$ is the exponent of the difference Δx. According to Lemma 1, this means that, for each $x \in [-\ell_{\max}, \ell_{\max}]$ the following relation must hold:

$$2^{e_{\frac{2}{\pi}} + e_x - p + 4} < 2^{e_{\Delta x}} \iff p > e_x - e_{\Delta x} + e_{\frac{2}{\pi}} + 4$$

The needed value of p can be obtained by finding the float $\hat{x} \in [-\ell_{\max}, \ell_{\max}]$ which maximizes the value of $e_{\hat{x}} - e_{\Delta\hat{x}}$. This task can be fulfilled by running the algorithm contained in [KM84].

It turns out that, for the IEEE 754 single precision floating point format, such $\hat{x}$ is $1.4\text{AC55C}_{16} \cdot 2^{21}$, with $|\Delta\hat{x}| = \left|\hat{x}\frac{2}{\pi} - \hat{k}\right| \simeq 1.\text{A1CCDF}_{16} \cdot 2^{-27}$. Running the algorithm in [KM84] with a "threshold" of $|\Delta x| \leq 2^{-25}$ suffices to show that for no other exponent $e_{\hat{x}'}$ there exists an $\hat{x}'$ such that $|\Delta\hat{x}'| < |\Delta\hat{x}|$. It follows that $p > 21 - 1 + 4 + 27 = 51$. Proceeding in the same way, the minimum precision needed for the Double Precision format is $p > 109$.

Since we established that for the Single format a precision of at least 52 bits is needed, IEEE 754 double precision numbers may be used to perform the multiplication $x \boxdot \left[\frac{2}{\pi}\right]$. The obtained result can be rounded to the nearest integer towards positive infinity using functions provided by the C and C++ standard math libraries. The cases of $x = -0.0$ and $x = +0.0$ must be handled separately: the values returned by this function must be 0 and 1 respectively, because of Definition 6.

Generation of lower $\frac{\pi}{2}$ multiple approximations A similar reasoning can be done to compute $\left[k\frac{\pi}{2}\right]_{\downarrow}$ for a given integer k, that is, the maximum floating point number $x \in \mathbb{F}$ such that $x < k\frac{\pi}{2}$. This can be achieved by multiplying k by the constant $\frac{\pi}{2}$ with a sufficient precision, and then rounding it to the target precision using the *roundTowardNegative* rounding direction.

The precision used for this multiplication must satisfy

$$p > e_k - e_{x - k\frac{\pi}{2}} + e_{\frac{\pi}{2}} + 4$$

according to Lemma 1. The value $\hat{x}$ that maximizes $e_x - e_{\Delta x}$ is the same as for the previous section. Since for each $x \in [-\ell_{\max}, \ell_{\max}]$ we have $|\Delta x| = \left|x\frac{2}{\pi} - k\right| \leq \left|x - k\frac{\pi}{2}\right|$ (where k is the integer nearest to $x\frac{2}{\pi}$), $e_{x-k\frac{\pi}{2}} \geq e_{\Delta x}$ holds. Also, because of the criteria we used to define $\ell_{\max}$ (i.e. the ulp in the interval should be sufficiently large, and greater than 1), k is always representable in our target floating-point format. Moreover, $e_k \leq e_x$, because $|k| \leq \left|\left[k\frac{\pi}{2}\right]_{\downarrow}\right| = |x|$. Consequently, $\hat{x}$ and $\hat{k}$ maximize also $e_k - e_{x-k\frac{\pi}{2}}$.

For the single precision format we can state

$$p > e_{\hat{x}} + e_{\frac{\pi}{2}} + 4 - e_{\Delta\hat{x}} = 21 + 0 + 4 + 27 = 52,$$

which imposes a required precision of at least 53 bits. This allows us to use the IEEE 754 double precision format to perform this multiplication. In order to round the result to the single precision format, the rounding direction can be set to *roundTowardNegative*. It is also possible to avoid switching rounding mode by adjusting the obtained result according to the rounding direction in use. If $k = 0$, the function should return $x = -0.0$. This case can be handled separately.

To obtain the result in the double precision format, $p > 110$ is required. The computation can be performed using a technique for the implementation of multiple precision arithmetic, such as the one described in [Dek71] with an extended precision format of at least 55 bits of precision, or the one given in [HLB01] with the double precision format. For formats requiring higher precision, see [MPT16].

5 Propagation Algorithms

Let $\mathbb{F}$ be any IEEE 754 floating-point format and let $\mathcal{S} \subseteq \wp(\mathbb{F})$ be a bounded meet-semilattice. A floating-point unary constraint over $\mathcal{A}$ is a formula of the form $\mathbf{x} \in S$ for $S \in \mathcal{S}$.

Let $f\colon \mathbb{F} \to \mathbb{F}$ be a function and consider a constraint of the form $\mathbf{y} = f(\mathbf{x})$ along with the unary constraints $\mathbf{x} \in S_x$ and $\mathbf{y} \in S_y$ with $S_x, S_y \in \mathcal{S}$.

Direct propagation amounts to compute a possibly refined set $S'_y \in \mathcal{S}$, such that

$$S'_y \subseteq S_y \wedge \forall x \in S_x : f(x) \in S_y \implies f(x) \in S'_y. \tag{1}$$

Of course this is always possible by taking $S'_y = S_y$, but the objective of the game is to compute a "small", possibly the smallest S'_y enjoying (1), compatibly with the availability of information on f and of computing resources. The smallest $S'_y \in \mathcal{S}$ that satisfies (1) is such that

$$\forall S''_y \in \mathcal{S} : S''_y \subset S'_y \implies \exists x \in S_x \,.\, f(x) \in S_y \setminus S''_y. \tag{2}$$

Inverse propagation for the same constraints, $\mathbf{y} = f(\mathbf{x})$, $\mathbf{x} \in S_x$ and $\mathbf{y} \in S_y$, is the computation of a possibly refined set for $\mathbf{x}$, S'_x, such that

$$S'_x \subseteq S_x \wedge \forall x \in S_x : f(x) \in S_y \implies x \in S'_x. \tag{3}$$

Again, taking $S'_x = S_x$ is always possible and sometimes unavoidable. The best we can hope for is to be able to determine the smallest such set, i.e., satisfying

$$\forall S''_x \in \mathcal{S} : S''_x \subset S'_x \implies \exists x \in S_x \setminus S''_x \,.\, f(x) \in S_y. \tag{4}$$

Satisfying predicates 2 and 4 would correspond to obtaining *domain consistency* [VSD98] on our constraint set. Unfortunately, this goal is often difficult

to reach, especially when the underlying variable domains are large. A less demanding approach to the problem is *interval consistency*: we associate an interval $[x_l, x_u]$ to variable $\mathbf{x}$ and an interval $[y_l, y_u]$ to $\mathbf{y}$.

If function f is isotonic, direct propagation can now be reduced to finding a new interval $[y'_l, y'_u]$ such that

$$y'_l \succcurlyeq y_l \wedge \forall x \in [x_l, x_u] : f(x) \succcurlyeq y_l \implies f(x) \succcurlyeq y'_l, \tag{5}$$
$$y'_u \preccurlyeq y_u \wedge \forall x \in [x_l, x_u] : f(x) \preccurlyeq y_l \implies f(x) \preccurlyeq y'_u. \tag{6}$$

Taking $y'_l = y_l$ and $y'_u = y_u$ would trivially satisfy these predicates. However, we aim to find an interval such that

$$\forall y''_l \in \mathbb{F} : y''_l \succ y'_l \implies \exists x \in [x_l, x_u] \,.\, y_l \preccurlyeq f(x) \prec y''_l, \tag{7}$$
$$\forall y''_u \in \mathbb{F} : y''_u \prec y'_u \implies \exists x \in [x_l, x_u] \,.\, y''_u \prec f(x) \preccurlyeq y_u. \tag{8}$$

Inverse propagation consists now in finding an interval $[x'_l, x'_u]$ such that

$$x'_l \succcurlyeq x_l \wedge \forall x \in [x_l, x_u] : f(x) \succcurlyeq y_l \implies x \succcurlyeq x'_l, \tag{9}$$
$$x'_u \preccurlyeq x_u \wedge \forall x \in [x_l, x_u] : f(x) \preccurlyeq y_u \implies x \preccurlyeq x'_u. \tag{10}$$

An optimal result would satisfy

$$\forall x''_l \in \mathbb{F} : x''_l \succ x'_l \implies \exists x \in [x_l, x''_l) \,.\, f(x) \succcurlyeq y_l, \tag{11}$$
$$\forall x''_u \in \mathbb{F} : x''_u \prec x'_u \implies \exists x \in (x''_u, x_u] \,.\, f(x) \preccurlyeq y_u. \tag{12}$$

Similar predicates can be conceived for antitonic functions.

A possible compromise between domain and interval consistency is the use of multi-intervals. It consists in splitting domains in multiple intervals, thus allowing more granularity. It is not difficult to extend the predicates given above to "multi-interval consistency".

Unfortunately, the functions we are concerned with are not isotonic, nor antitonic, because of the presence of glitches. However, we devised algorithms that, given the implementation of a quasi-monotonic library function $f\colon \mathbb{F} \to \mathbb{F}$, an interval $[x_l, x_u]$ for $\mathbf{x}$ and an interval $[y_l, y_u]$ for $\mathbf{y}$, can compute refined bounds for both intervals, satisfying the correctness predicates defined above and, in particular cases, even the optimality predicates. These algorithms, that will be presented in the following sections, exploit some simple data describing the glitches of a specific function in order to overcome the issues generated by its quasi-monotonicity. Such data consists in safe approximations n_{g}, d_{M} and w_{M} of, respectively, the total number of glitches n_{g}^f, their maximal depth d_{M}^f and width w_{M}^f. Moreover, a safe approximation α of where the first glitch starts (α^f) and, analogously, a safe approximation ω of where the last glitch ends (ω^f) inside the function's domain are needed.

5.1 Direct Propagation

Given an interval $[x_l, x_u]$ for $\mathbf{x}$ and a function f, finding a refined interval $[y_l, y_u]$ for $\mathbf{y}$ that satisfies constraint $\mathbf{y} = f(\mathbf{x})$ would be trivial if f was monotonic:

computing $[y'_l, y'_u] \equiv [f(x_l), f(x_u)]$ would suffice. However, the possible presence of glitches in quasi-monotonic functions requires some more care. In particular, it raises two main issues:

- there might be glitches in $[x_l, x_u]$ in which the value of the function is lower than $f(x_l)$;
- x_u might be inside a glitch, and there might be values of x outside it such that $f(x) \succ f(x_u)$.

If $[x_l, x_u]$ and $[\alpha, \omega]$ do not intersect, the simple approach for monotonic functions can be used. Otherwise, the said problems can be tackled in the following ways, exploiting some minimal information about the glitches of the function.

y_l : If $x_l \in [\alpha, \omega]$, the worst-case scenario is that there is a glitch starting right after x_l, where the graph of f goes lower than $f(x_l)$. However, the said glitch surely cannot be deeper than $f(x_l)^{-^{d_\mathrm{M}}}$: we can take this value as y_l, the lower bound of the refined interval.
If x_l is not in the "glitch-area", then we can consider the value of f in α: surely no glitch in $[\alpha, x_u]$ can go lower than $f(\alpha)^{-^{d_\mathrm{M}}}$. In this case, we can set $y_l = \min\{f(x_l), f(\alpha)^{-^{d_\mathrm{M}}}\}$.

y_u : If $x_u \notin [\alpha, \omega]$, then it cannot be in a glitch, and $f(x_u)$ is the highest value of f in $[x_l, x_u]$. Otherwise, it may be in a glitch, but surely not deeper than d_M: the actual maximum value of the function, right outside the glitch, cannot be higher than $f(x_u)^{+^{d_\mathrm{M}}}$. y_u can be set to this value.

If the maximum and minimum values taken by the function in its whole domain are also known, they can be compared to the obtained values of y_l and y_u in order not to return values outside the range of f.

5.2 Indirect Propagation

For inverse propagation, i.e. the process of inferring a new interval $[x_l, x_u]$ for $\mathbf{x}$ starting from the interval $[y_l, y_u]$ for $\mathbf{y}$, we decided to split the problem, by separately finding a lower bound x'_l and an upper bound x'_u for values of x that satisfy the equation $\mathbf{y} = f(\mathbf{x})$. We therefore designed two algorithms, `lower_bound` and `upper_bound`, described in the following sections, capable of carrying out the tasks suggested by their names, but for the equation $y = f(\mathbf{x})$, considering a single value of y.

Function `lower_bound`, reported in Algorithm 1, returns a value l that satisfies one of the following predicates:

$$
\begin{aligned}
p_0(y, x_l, x_u, l) &\equiv \forall x \in [x_l, x_u] : y \succ f(x), \\
p_1(y, x_l, x_u, l) &\equiv \forall x \in [x_l, l] : y \prec f(x), \\
p_2(y, x_l, x_u, l) &\equiv \forall x \in [x_l, l] : y \succ f(x), \\
p_3(y, x_l, x_u, l) &\equiv f(l) \prec y \prec f(l^+) \wedge \forall x \in [x_l, l) : y \succ f(x), \\
p_4(y, x_l, x_u, l) &\equiv y = f(l) \wedge \forall x \in [x_l, l) : y \succ f(x).
\end{aligned}
$$

When condition $p_0(y, x_l, x_u, l)$ holds, the equation $y = f(\mathbf{x})$ has no solution. Moreover, even when $p_1(y, x_l, x_u, l)$ holds with $l = x_u$, the equation $y = f(\mathbf{x})$ has no solution. When $p_1(y, x_l, x_u, l)$ holds with $l \prec x_u$ or $p_2(y, x_l, x_u, l)$ holds, the equation $y = f(\mathbf{x})$ may have a solution and choosing $x'_l = l^+$ is correct. When $p_3(y, x_l, x_u, l)$ holds, we identified the leftmost point in $[x_l, x_u]$ where f crosses y without touching it and we can set $x'_l = l^+$. Finally, when $p_4(y, x_l, x_u, l)$ holds we identified the leftmost solution $x = l$ of $y = f(\mathbf{x})$ and we can set $x'_l = l$.

The upper bound is instead computed by function `upper_bound`, reported in Algorithm 4. It terminates returning a value u that satisfies the predicates listed below:

$$
\begin{aligned}
&p_5(y, x_l, x_u, u) \equiv \forall x \in [x_l, x_u] : y \prec f(x),\\
&p_6(y, x_l, x_u, u) \equiv \forall x \in [u, x_u] : y \succ f(x),\\
&p_7(y, x_l, x_u, u) \equiv \forall x \in [u, x_u] : y \prec f(x),\\
&p_8(y, x_l, x_u, u) \equiv f(u^-) \prec y \prec f(u) \wedge \forall x \in (u, x_u] : y \prec f(x),\\
&p_9(y, x_l, x_u, u) \equiv y = f(u) \wedge \forall x \in (u, x_u] : y \prec f(x).
\end{aligned}
$$

These predicates are the counterparts of those for `lower_bound`, previously presented. If one of predicates $p_5(y, x_l, x_u, u)$ and $p_6(y, x_l, x_u, u)$ holds, the latter only with $u = x_l$, then equation $y = f(\mathbf{x})$ has no solution in $[x_l, x_u]$. If $p_6(y, x_l, x_u, u)$ holds with $u \prec x_u$, or $p_7(y, x_l, x_u, u)$ holds with any $u \in [x_l, x_u]$, then equation $y = f(\mathbf{x})$ has surely no solution in interval $[u, x_u]$, but it might have a solution somewhere in $[x_l, u)$. Setting $x'_u = u^-$ is therefore correct. If $p_8(y, x_l, x_u, u)$ holds, then the algorithm identified the rightmost point in $[x_l, x_u]$ where the graph of f crosses y without touching it. Setting $x'_u = u^-$ is correct. Finally, when $p_9(y, x_l, x_u, u)$ holds, the algorithm found the rightmost solution $x = u$ of the equation $y = f(\mathbf{x})$ and setting $x'_u = u$ is correct.

If function f is quasi-isotonic, a lower bound x'_l for $\mathbf{x}$ satisfying equation $\mathbf{y} = f(\mathbf{x})$ can be found by calling `lower_bound` with $y = y_l$, and `upper_bound` with $y = y_u$. The results of the two algorithms are then combined depending on which of predicates p_i, $i \in \{0, \ldots, 9\}$ they satisfy.

- If one of predicates p_0 or p_5 holds, then we are sure that there is no solution, because the entire graph of the function is either below or above the interval for $\mathbf{y}$. Note that p_0 implies p_6 and p_5 implies p_1.
- If p_1 holds, then we must set $x'_l = x_l$ even if $l \succ x_l$. In fact, although we are sure the graph of the function is entirely above y_l, there might be a value y'_l such that $y_l \prec y'_l \preccurlyeq y_u$ that satisfies $y'_l = f(\mathbf{x})$ for some $x \in [x_l, l]$. Excluding such interval would therefore undermine the correctness of the propagation process. This, of course, unless either $\mathbf{y}$ is a singleton, or also p_5 holds: in these cases we can correctly conclude that there are no solutions. The same reasoning can be done for the upper bound: unless $\mathbf{y}$ is a singleton or p_0 holds, x'_u must be set to x_u whenever p_6 holds.
- If p_3 and p_8 hold together and $l \succ u$, then we are sure that there is no solution for equation $\mathbf{y} = f(\mathbf{x})$.

– For all other combinations of predicates, the new values x'_l and x'_u can be set as described below the definitions of the predicates.

The same algorithms can be used for the quasi-antitonic functions, by exploiting the fact that, if f is quasi-antitonic, then $-f$ is quasi-isotonic. Both algorithms can therefore be called with $f' = -f$, $f^{\mathrm{i}\prime} = f \circ (-\mathrm{id})$, and $-y_u$ instead of y_l for `lower_bound`, and $-y_l$ in place of y_u for `upper_bound`. When they terminate, $p_i(-y_u, x_l, x_u, l)$ and $p_j(-y_l, x_l, x_u, u)$ with $i \in \{0, \ldots, 4\}$ and $j \in \{5, \ldots, 9\}$ hold with $-f$. Note that, since $-y \prec -f(x) \iff y \succ f(x)$ and $-y \succ -f(x) \iff y \prec f(x)$, the same predicates do not hold on the original function f. However, because y_l and y_u are switched, the same discussions regarding the treatment of the new bounds x'_l and x'_u depending on the post-condition predicates we did for the isotonic case can be done, obtaining the same guidelines with the same values of i and j.

Computation of the Lower Bound for $y = f(\mathbf{x})$ Given a value for y and equation $y = f(\mathbf{x})$, Algorithm 1 computes a correct lower bound that (possibly) allows us to refine the interval of $\mathbf{x}$. The preconditions are listed in the **Require** statement and demand a quasi-isotonic function $f\colon \mathbb{F} \to \mathbb{F}$, a value $y \in \mathbb{F}$, and an interval $[x_l, x_u]$ for $\mathbf{x}$ to be refined. In order to be as precise as possible, the algorithm needs safe approximations of the glitch data listed previously in this section. It also uses an inverse function $f^{\mathrm{i}}\colon \mathbb{F} \to \mathbb{F}$, if available. Finally, the parameter $t \in \mathbb{N}$ is the maximum length of the linear searches that the algorithm performs in some cases, and $s \in \mathbb{N}$ is the maximum number of times the function `logsearch_lb` in Algorithm 3 can return an interval too wide for ensuring the logarithmic complexity of the dichotomic process.

The algorithm ends guaranteeing the post-condition expressed in the **Ensure** statement, where the predicates $p_r(y, x_l, x_u, l)$ for $r \in \{0, \ldots, 4\}$ are those described previously. The objective of the algorithm is to determine r and l by performing a number of calls to the concerned library function bounded by a small constant k (e.g., $k = 3$), times the logarithm of the cardinality of interval $[x_l, x_u]$.

First, Algorithm 1 calls function `init`, that takes a value y, interval $[x_l, x_u]$ for $\mathbf{x}$, and an inverse function $f^{\mathrm{i}}\colon \mathbb{F} \to \mathbb{F}$, if available. Its purpose is to simply return a point inside interval $[x_l, x_u]$. If the inverse function f^{i} is available and $x_l \preccurlyeq f^{\mathrm{i}}(y) \preccurlyeq x_u$, then `init` returns $f^{\mathrm{i}}(y)$; it returns the middle point between x_l and x_u otherwise.

Next, function `gallop_lb` is called. Its goal is to find values lo and hi that satisfy the precondition of algorithm `bisect_lb` and are suitable for the bisection process, i.e. such that $x_l \preccurlyeq \mathrm{lo} \preccurlyeq \mathrm{hi} \preccurlyeq x_u$, $(x_l \prec \mathrm{lo}) \implies y \succ^{d_{\mathrm{M}}} f(\mathrm{lo})$ and $x_u \succ \mathrm{hi} \implies f(\mathrm{hi}) \succcurlyeq y$. Therefore, `gallop_lb` starts with $\mathrm{hi} = i$ and increases it (e.g. by multiplying it by 2, if positive) until a new hi such that $\mathrm{hi} \prec x_u$ and $f(\mathrm{hi}) \succcurlyeq y$ is found. If such hi cannot be found in this way, it sets $\mathrm{hi} = x_u$. Analogously, function `gallop_lb` starts with $\mathrm{lo} = i$ and then decreases lo (e.g. by dividing it by 2, if positive) until a new lo such that $x_l \prec \mathrm{lo}$ and $y \succ^{d_{\mathrm{M}}} f(\mathrm{lo})$ is found. If such lo cannot be found in this way, it sets $\mathrm{lo} = x_l$.

Algorithm 1 Inverse propagation: `lower_bound`$(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, f^\mathrm{i}, s, t)$

Require: $f\colon \mathbb{F} \to \mathbb{F}$, $y \in \mathbb{F}$, $[x_l, x_u] \in \mathcal{I}_\mathbb{F}$, $n_\mathrm{g} \geq n_\mathrm{g}^f$, $d_\mathrm{M} \geq d_\mathrm{M}^f$, $w_\mathrm{M} \geq w_\mathrm{M}^f$, $\alpha \preccurlyeq \alpha^f$, $\omega \succcurlyeq \omega^f$, $n_\mathrm{g} > 0 \implies (x_l \preccurlyeq \alpha \preccurlyeq \omega \preccurlyeq x_u)$, $f^\mathrm{i}\colon \mathbb{F} \to \mathbb{F}$, $s, t \in \mathbb{N}$.
Ensure: Correctness: $l \in \mathbb{F}$, $r \in \{0,1,2,3,4\} \implies p_r(y, x_l, x_u, l)$
Precision: $\Big(f(x_l) \preccurlyeq y \preccurlyeq f(x_u) \wedge \big(n_\mathrm{g} = 0 \vee w_\mathrm{M} < t \vee (n_\mathrm{g} = 1 \wedge \alpha = \alpha^f)\big)\Big) \implies r \in \{3,4\}$

$i := \texttt{init}(y, [x_l, x_u], f^\mathrm{i})$
▷ $x_l \preccurlyeq i \preccurlyeq x_u$
$(\mathrm{lo}, \mathrm{hi}) := \texttt{gallop_lb}(f, y, [x_l, x_u], d_\mathrm{M}, i)$
▷ $(x_l \preccurlyeq \mathrm{lo} \preccurlyeq \mathrm{hi} \preccurlyeq x_u) \wedge (x_l \prec \mathrm{lo} \implies y \succ^{d_\mathrm{M}} f(\mathrm{lo})) \wedge (x_u \succ \mathrm{hi} \implies f(\mathrm{hi}) \succcurlyeq y)$
if $f(\mathrm{lo}) \succ y$ **then**
  if $n_\mathrm{g} = 0 \vee f(\alpha) \succ^{d_\mathrm{M}} y$ **then**
    $l = x_u$; $r := 1$; **return**
  else
    $l := \alpha$; $r := 1$; **return**
  end if
else if $f(\mathrm{lo}) = y$ **then**
  $l := \mathrm{lo}$; $r := 4$; **return**
end if
if $f(\mathrm{hi}) \prec y$ **then**
  $(r, l, \mathrm{hi}) := \texttt{findhi_lb}(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, t)$
  if $r \in \{0,2\}$ **then return**
  end if
end if
$\mathrm{lo} := \texttt{bisect_lb}(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, s, t, \mathrm{lo}, \mathrm{hi})$
while $f(\mathrm{lo}^+) \prec y \wedge t > 0$ **do**
  $\mathrm{lo} := \mathrm{lo}^+$;
  $t := t - 1$
end while
if $f(\mathrm{lo}^+) \succ y$ **then**
  $l := \mathrm{lo}$; $r := 3$
else if $f(\mathrm{lo}^+) = y$ **then**
  $l := \mathrm{lo}^+$; $r := 4$
else
  $l := \mathrm{lo}$; $r := 2$
end if

The latter case, in which $f(\mathrm{lo}) \succ y$, is handled at lines 5-10. We need to determine if really $[x_l, x_u]$ does not contain any solution for y, that is, if $\forall x \in [x_l, x_u] : y \succ f(x)$ holds. If we are not sure because glitches are too deep, a suboptimal value for l is returned, and the algorithm terminates. If $f(\mathrm{lo}) = y$, at line 11 the algorithm returns with the exact solution for the lower bound.

Algorithm 2 Inverse propagation: `findhi_lb`$(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, t)$

Require: $f\colon \mathbb{F} \to \mathbb{F}$, $y \in \mathbb{F}$, $[x_l, x_u] \in \mathcal{I}_\mathbb{F}$, $n_\mathrm{g} \geq n_\mathrm{g}^f$, $d_\mathrm{M} \geq d_\mathrm{M}^f$, $w_\mathrm{M} \geq w_\mathrm{M}^f$, $\alpha \preccurlyeq \alpha^f$, $\omega \succcurlyeq \omega^f$, $n_\mathrm{g} > 0 \implies (x_l \preccurlyeq \alpha \preccurlyeq \omega \preccurlyeq x_u)$, $t \in \mathbb{N}$, $f(x_u) \prec y$.
Ensure: $l \in \mathbb{F}$, $r \in \{0, 2\} \implies p_r(y, x_l, x_u, l)$, $r = 1 \implies \big(\mathrm{hi} \in [x_l, x_u] \wedge f(\mathrm{hi}) \succcurlyeq y\big)$.

```
 1: l = x_l;
 2: if n_g = 0 ∨ x_u ≻ ω ∨ y ≻^{d_M} f(x_u) then
 3:     r := 0
 4: else if n_g = 1 ∧ (w_M > t ∨ f(α⁺) ≺ f(α)) then
 5:     if y ≻ f(α) then
 6:         if f(α⁺) ≺ f(α) then
 7:             r := 0
 8:         else
 9:             l := α; r := 2
10:         end if
11:     else
12:         hi := α; r := 1
13:     end if
14: else
15:     (b, hi) := linsearch_geq(f, y, [x_l, x_u], w_M, t)
16:             ▷ (b = 1 ∧ hi ∈ [x_l, x_u] ∧ f(hi) ≽ y) ∨ (b = 0 ∧ ∀x ∈ [x̂, x_u] : f(x) ≺ y)
17:                                   ▷ where v = min{t, w_M} and x̂ = max{x_l, x_u^{-v}}
18:     if b = 1 then
19:         r := 1
20:     else if t ≥ w_M then
21:         r := 0
22:     else
23:         l := x_l; r := 2
24:     end if
25: end if
```

On line 15, function `findhi_lb`, defined by Algorithm 2, handles the special case where `gallop_lb` ended up with $\mathrm{hi} = x_u$ and $f(x_u) \prec y$. This might happen because for no $x \in [x_l, x_u]$ we have $f(x) \succcurlyeq y$, or because such an x does exist, but x_u is in a glitch. The purpose of `findhi_lb` is to quickly discriminate between these two cases and find a value for hi suitable for bisection, when possible. In this case function `findhi_lb` returns $r = 1$.

- If x_u might be in a (wider than t) glitch, for the sake of efficiency we cannnot exhaustively look for good candidates to apply bisection. However, we might still be able to set $r = 0$, to signify that $\forall x \in [x_l, x_u] : f(x) \prec y$, or to set

$r = 1$ and hi to an appropriate value for bisection. If we do not have enough information to chose one of these options, we may just set $r = 2$ and l to a valid (possibly suboptimal) lower bound.

- Otherwise, `linsearch_geq`$(f, y, [x_l, x_u], w_{\mathrm{M}}, t)$ performs a backwards, float-by-float, search for at most the minimum between t and w_{M} steps, starting from $\mathrm{hi} = x_u$, and looking for the first value for hi such that $f(\mathrm{hi}) \succcurlyeq y$. The search stops in two cases, which are discriminated by the value given to variable b.
  - $b = 0$: within the allowed t steps of the search, no suitable value for hi was found. Still, if t was high enough to cover the glitch, we set $r = 0$ since we are sure that $\forall x \in [x_l, x_u] : f(x) \prec y$. Otherwise we set $r = 2$ and we content ourselves with the trivial lower bound x_l.
  - $b = 1$: an appropriate value hi for bisection has been found.

At the end of line 17, we know that function `findhi_lb` returned $r = 1$ and, therefore, a new value for hi was returned such that $x_l \preccurlyeq \mathrm{hi} \preccurlyeq x_u$ and $f(\mathrm{hi}) \succcurlyeq y$.

Summarising, before the invocation of function `bisect_lb` at line 19, we are sure that $f(\mathrm{lo}) \prec y \preccurlyeq f(\mathrm{hi})$ and, therefore, $\mathrm{lo} \neq \mathrm{hi}$. Function `bisect_lb`, defined by Algorithm 3, performs a dichotomic search to refine the interval $[\mathrm{lo}, \mathrm{hi}]$ when f is a quasi-isotonic function. Each iteration of the **while** loop on line 1 uses function `split_point` to pick the middle point mid of interval $[\mathrm{lo}, \mathrm{hi}]$, and compares the value of $f(\mathrm{mid})$ with y. When $f(\mathrm{mid}) \succcurlyeq y$, the value of hi can be correctly updated with the value of mid. The critical case is when $f(\mathrm{mid}) \prec y$. In this case, function `bisect_lb` has to further discriminate whether lo can be correctly updated with the value of mid or other refinements of the interval $[\mathrm{lo}, \mathrm{hi}]$ are possible. In particular:

- if mid is surely not inside a glitch, the value of lo can correctly be updated with the value of mid;
- if mid *may be inside* the unique glitch that either is bigger than t or starts exactly in α (that is $f(\alpha^+) \prec f(\alpha)$), the algorithm needs to test the value of f in α and ω in order to determine if and how the interval $[\mathrm{lo}, \mathrm{hi}]$ can be correctly refined;
- if mid *may be inside* a glitch narrower than t, function `findfmax` finds the value b inside interval $[\max\{\mathrm{lo}, \mathrm{mid}^{-w_{\mathrm{M}}}\}, \mathrm{mid}]$ for which value $f(b)$ is the highest. Such value is then used to refine interval $[\mathrm{lo}, \mathrm{hi}]$;
- if mid *may be inside* a glitch wider than t, for complexity reasons Algorithm 3 refrains from iterating the float-by-float search of function `findfmax` and calls `logsearch_lb` that finds, if it exists, a value $z \in [\mathrm{lo}, \mathrm{mid}]$ such that $f(z) \prec^{d_{\mathrm{M}}} y$. If such z was found, it is used to refine interval $[\mathrm{lo}, \mathrm{hi}]$.

In more detail, the code of `logsearch_lb` performs a logarithmic search for finding, if it exists, a value $z \in [\mathrm{lo}, \mathrm{mid}]$ such that $f(z) \prec^{d_{\mathrm{M}}} y$. Its argument $s \in \mathbb{N}$ constitutes (for efficiency reasons) an upper bound to the number of times the code of `logsearch_lb` can return an excessively wide interval as refinement of $[x_l, x_u]$. If such s is not reached, `logsearch_lb` starts with $z = \mathrm{mid}$ and decreases

Algorithm 3 Inverse propagation: `bisect_lb`$(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, s, t,$ lo, hi)

Require: $x_l \preccurlyeq \mathrm{lo} \prec \mathrm{hi} \preccurlyeq x_u$. $f(\mathrm{lo}) \prec y \preccurlyeq f(\mathrm{hi})$, $\forall x \in [x_l, \mathrm{lo}] : f(x) \prec y$ $f\colon \mathbb{F} \to \mathbb{F}$, $y \in \mathbb{F}$, $[x_l, x_u] \in \mathcal{I}_\mathbb{F}$, $n_\mathrm{g} \geq n_\mathrm{g}^f$, $d_\mathrm{M} \geq d_\mathrm{M}^f$, $w_\mathrm{M} \geq w_\mathrm{M}^f$, $\alpha \preccurlyeq \alpha^f$, $\omega \succcurlyeq \omega^f$, $n_\mathrm{g} > 0 \implies (x_l \preccurlyeq \alpha \preccurlyeq \omega \preccurlyeq x_u)$, $s, t \in \mathbb{N}$.

Ensure: Correctness: $x_l \preccurlyeq \mathrm{lo} \prec \mathrm{hi} \preccurlyeq x_u$, $f(\mathrm{lo}) \prec y \preccurlyeq f(\mathrm{hi})$, $\forall x \in [x_l, \mathrm{lo}] : f(x) \prec y$,
Precision: $\big(n_\mathrm{g} = 0 \vee w_\mathrm{M} < t \vee (n_\mathrm{g} = 1 \wedge \alpha = \alpha^f)\big) \implies f(\mathrm{lo}^+) \succcurlyeq y$

while $\mathrm{hi} \succ^1 \mathrm{lo}$ **do**
  $\mathrm{mid} :=$ `split_point`$(\mathrm{lo}, \mathrm{hi})$
  $\triangleright\ \exists m, m' > 0,\ |m - m'| \leq 1,\ \mathrm{mid} = \mathrm{hi}^{-m} = \mathrm{lo}^{+m'}$
  if $y \preccurlyeq f(\mathrm{mid})$ **then**
    $\mathrm{hi} := \mathrm{mid}$
  else if $n_\mathrm{g} = 0 \vee \mathrm{mid} \preccurlyeq \alpha \vee \mathrm{mid} \succcurlyeq \omega \vee y \succ^{d_\mathrm{M}} f(\mathrm{mid})$ **then**
    $\mathrm{lo} := \mathrm{mid}$
  else if $n_\mathrm{g} = 1 \wedge \big(w_\mathrm{M} > t \vee f(\alpha^+) \prec f(\alpha)\big)$ **then**
    if $f(\omega) \succcurlyeq y$ **then**
      if $f(\alpha) \succcurlyeq y$ **then**
        $\mathrm{hi} := \alpha$
      else if $f(\alpha^+) \prec f(\alpha)$ **then**
        $\mathrm{lo} := \mathrm{mid}$
      else if $\mathrm{lo} \prec \alpha$ **then**
        $\mathrm{lo} := \alpha$;
      else
        break
      end if
    else $\mathrm{lo} := \omega$
    end if
  else if $w_\mathrm{M} \leq t$ **then**
    $b :=$ `findfmax`$(f, w_\mathrm{M}, \mathrm{lo}, \mathrm{mid})$
    $\triangleright\ b \in [\max\{\mathrm{lo}, \mathrm{mid}^{-w_\mathrm{M}}\}, \mathrm{mid}] \wedge \forall x \in [\max\{\mathrm{lo}, \mathrm{mid}^{-w_\mathrm{M}}\}, \mathrm{mid}],\ f(x) \preccurlyeq f(b)$
    if $f(b) \succcurlyeq y$ **then**
      $\mathrm{hi} := b$
    else
      $\mathrm{lo} := \mathrm{mid}$
    end if
  else
    $z :=$ `logsearch_lb`$(f, d_\mathrm{M}, \mathrm{lo}, \mathrm{mid}, y, s)$
    $\triangleright\ z \in [\mathrm{lo}, \mathrm{mid}] \wedge \big((\mathrm{lo} \prec z) \implies f(z) \prec^{d_\mathrm{M}} y\big)$
    if $\mathrm{lo} \prec z$ **then**
      $\mathrm{lo} := z$
    else
      break
    end if
  end if
end while

it (e.g. by dividing it by 2, if positive) until a new z such that $f(z) \prec^{d_M} y$ is found. Otherwise, z is set to lo.

After `bisect_lb` has terminated, its post-condition ensures that $x_l \preccurlyeq \text{lo} \prec \text{hi} \preccurlyeq x_u$, $f(\text{lo}) \prec y \preccurlyeq f(\text{hi})$ and $\forall x \in [x_l, \text{lo}] : f(x) \prec y$ hold. Then, a **while** loop is entered. This loop performs a float-by-float search (for no more than t iterations) to approach the exact solution of $y = f(\mathbf{x})$. Afterwards, the invariant of the loop $\forall x \in [x_l, \text{lo}] : f(x) \prec y$ holds. At line 24 and line 28 we test if an optimal solution of $y = f(\mathbf{x})$ was found. If the body of the **else** statement of line 28 is reached, it means that the **while** loop of line 20 terminated because t reached 0. In this case, l is set to lo, a suboptimal solution of $y = f(\mathbf{x})$.

The code of the above mentioned functions `init`, `gallop_lb`, `linsearch_geq`, `findfmax` and `logsearch_lb` is not shown because it is straightforward. The formal proofs of the correctness of Algorithms 1, 2 and 3 can be found in Theorem 1 and Lemmas 2 and 3 of Appendix B.1, respectively. These results also show that when the function f is isotonic or it has glitches narrower than t, Algorithm 1 finds a precise solution, that is, it returns either $r = 3$ or $r = 4$.

The complexity of the code of function `lower_bound` in Algorithm 1, measured as the number of calls to function f, in the worst case is $k \log_2(\#[x_l, x_u]) + k + c$ for small constants k and c. For example, if f has more than one glitch and $w_M \leq t$, k and c are strictly related to w_M. Theorems 2 and 3 of Appendix B.1 prove the exact complexity of Algorithm 1 for some special cases.

Computation of the Upper Bound for $y = f(\mathbf{x})$ Algorithm 4 tries to refine the interval for $\mathbf{x}$ by finding a correct upper bound. The preconditions for this algorithm are listed in its **Require** statement, and they are the same as those of Algorithm 1.

The algorithm terminates satisfying the post-condition presented in its **Ensure** statement. The predicates p_i that appear in the statement are those defined at the beginning of Section 5.2.

The way algorithm `upper_bound` operates is substantially similar to the one of `lower_bound`. First, function `init` tries to guess initial values for lo and hi by using the rough inverse f^{i}. Then, `gallop_ub` tries to quickly find values for lo and hi suitable for the call to `bisect_ub`. In particular, $f(\text{lo}) \preccurlyeq y \prec f(\text{hi})$ and $\forall x \in [\text{hi}, x_u] : f(x) \succ y$ must hold. If this is not the case, the **if** statement of the subsequent lines handle the situation.

If $f(\text{hi}) \prec y$, we must discern the case when there is no solution from the case when x_u is in a glitch, and there is actually a value of x such that $f(x)$ reaches y somewhere outside of the glitch. Function `findhi_ub` is called for this purpose. It operates mostly like `findhi_lb`: if we can be sure that x_u is not in a glitch, either because there are no glitches, or because it is outside of $[\alpha, \omega]$, or if anyways no glitch could be deep enough for the function to reach y outside of it, we can claim that the function is lower than y for the while interval. Otherwise, if there is only one glitch and we know its position, some reasonings can be done to conclude whether there is a solution or not. If none of these conditions apply, a float-by-float search is performed in order to try to reach the beginning

Algorithm 4 Inverse propagation: `upper_bound`$(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, f^\mathrm{i}, s, t)$

Require: $f\colon \mathbb{F} \to \mathbb{F}$, $y \in \mathbb{F}$, $[x_l, x_u] \in \mathcal{I}_\mathbb{F}$, $n_\mathrm{g} \geq n_\mathrm{g}^f$, $d_\mathrm{M} \geq d_\mathrm{M}^f$, $w_\mathrm{M} \geq w_\mathrm{M}^f$, $\alpha \preccurlyeq \alpha^f$, $\omega \succcurlyeq \omega^f$, $n_\mathrm{g} > 0 \implies (x_l \preccurlyeq \alpha \preccurlyeq \omega \preccurlyeq x_u)$, $f^\mathrm{i}\colon \mathbb{F} \to \mathbb{F}$, $s, t \in \mathbb{N}$.

Ensure: Correctness: $u \in \mathbb{F}$, $r \in \{5,6,7,8,9\} \implies p_r(y, x_l, x_u, u)$

Precision: $\Big(f(x_l) \preccurlyeq y \preccurlyeq f(x_u)$

$\wedge\Big(n_\mathrm{g} = 0 \vee \big(n_\mathrm{g} = 1 \wedge w_\mathrm{M} < t \wedge \big(\alpha = \alpha^f \vee \omega = \omega^f \vee (\alpha \prec^k \omega \wedge k \leq t)\big)\big)\Big)\Big)$

$\implies r \in \{8,9\}$

1: $i := \texttt{init}(y, [x_l, x_u], f^\mathrm{i})$ ▷ $x_l \preccurlyeq i \preccurlyeq x_u$
2: $(\mathrm{lo}, \mathrm{hi}) := \texttt{gallop_ub}(f, y, [x_l, x_u], d_\mathrm{M}, i)$
3: ▷ $(x_l \preccurlyeq \mathrm{lo} \preccurlyeq \mathrm{hi} \preccurlyeq x_u) \wedge (x_l \prec \mathrm{lo} \implies y \succcurlyeq f(\mathrm{lo})) \wedge (x_u \succ \mathrm{hi} \implies f(\mathrm{hi}) \succ^{d_\mathrm{M}} y)$
4: **if** $f(\mathrm{hi}) \prec y$ **then**
5: $u := \texttt{findhi_ub}(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, t)$
6: $r := 6$; **return**
7: **else if** $f(\mathrm{hi}) = y$ **then**
8: $u := \mathrm{hi}$; $r := 9$; **return**
9: **end if**
10: **if** $f(\mathrm{lo}) \succ y$ **then**
11: **if** $n_\mathrm{g} = 0 \vee f(\alpha) \succ^{d_\mathrm{M}} y$ **then**
12: $r := 5$; **return**
13: **else**
14: $(b, z) := \texttt{check_glitch}(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, t, \mathrm{lo}, \mathrm{lo}, \mathrm{hi})$
15: **if** $b = 0$ **then**
16: $r := 5$; **return**
17: **else if** $b = 1$ **then**
18: **if** $f(z) = y$ **then**
19: $u = z$; $r = 9$; **return**
20: **else**
21: $u = z^+$; $r = 8$; **return**
22: **end if**
23: **else**
24: $u := \min\{\mathrm{hi}, \omega\}$; $r := 7$; **return**
25: **end if**
26: **end if**
27: **end if**
28: $\mathrm{hi} := \texttt{bisect_ub}(f, y, n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, n_\mathrm{g}, s, t, \mathrm{lo}, \mathrm{hi})$
29: **while** $f(\mathrm{hi}^-) \succ y \wedge t > 0$ **do**
30: $\mathrm{hi} := \mathrm{hi}^-$
31: $t := t - 1$
32: **end while**
33: **if** $f(\mathrm{hi}^-) \prec y$ **then**
34: $u := \mathrm{hi}$; $r := 8$
35: **else if** $f(\mathrm{hi}^-) = y$ **then**
36: $u := \mathrm{hi}^-$; $r := 9$
37: **else**
38: $u := \mathrm{hi}$; $r := 7$
39: **end if**

Algorithm 5 Inverse propagation: `findhi_ub`$(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, t)$

Require: $f\colon \mathbb{F} \to \mathbb{F}$, $y \in \mathbb{F}$, $[x_l, x_u] \in \mathcal{I}_\mathbb{F}$, $n_\mathrm{g} \geq n_\mathrm{g}^f$, $d_\mathrm{M} \geq d_\mathrm{M}^f$, $w_\mathrm{M} \geq w_\mathrm{M}^f$, $\alpha \preccurlyeq \alpha^f$, $\omega \succcurlyeq \omega^f$, $n_\mathrm{g} > 0 \implies (x_l \preccurlyeq \alpha \preccurlyeq \omega \preccurlyeq x_u)$, $t \in \mathbb{N}$, $f(x_u) \prec y$.
Ensure: $u \in \mathbb{F}$, $p_6(y, x_l, x_u, u)$.

if $n_\mathrm{g} = 0 \vee x_u \succ \omega \vee y \succ^{d_\mathrm{M}} f(x_u)$ **then**
  $u := x_l$
else if $n_\mathrm{g} = 1 \wedge \big(w_\mathrm{M} > t \vee (f(\alpha^+) \prec f(\alpha) \wedge y \succcurlyeq f(\alpha))\big)$ **then**
  if $y \prec f(\alpha) \vee f(\alpha^+) \succcurlyeq f(\alpha)$ **then**
    $u := x_u$
  else if $y = f(\alpha)$ **then**
    $u := \alpha^+$
  else
    $u := x_l$
  end if
else
  $(b, \mathrm{hi}, \hat{x}) :=$ `linsearch_geq`$(f, y, [x_l, x_u], w_\mathrm{M}, t)$
    ▷ $\big(b = 1 \wedge \mathrm{hi} \in [x_l, x_u] \wedge f(\mathrm{hi}) \succcurlyeq y \wedge \forall x \in [\mathrm{hi}^+, x_u] : f(x) \prec y\big)$
    ▷ $\vee \big(b = 0 \wedge \forall x \in [\hat{x}, x_u] : f(x) \prec y\big)$
    ▷ where $v = \min\{t, w_\mathrm{M}\}$ and $\hat{x} = \max\{x_l, {x_u}^{-v}\}$
  if $b = 1$ **then**
    $u := \mathrm{hi}^+$
  else if $t \geq w_\mathrm{M}$ **then**
    $u := x_l$
  else
    $u := \hat{x}$
  end if
end if

of the glitch where x_u is, if any. Since predicate p_6 always holds when function `findhi_ub` terminates, it cannot find a value of hi suitable for bisection, and it is only useful when the interval for $\mathbf{y}$ is a singleton.

Algorithm 6 Inverse propagation: `check_glitch`$(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, t, \mathrm{lo}, \mathrm{m}, \mathrm{hi})$

Require: $x_l \preccurlyeq \mathrm{lo} \preccurlyeq \mathrm{m} \preccurlyeq \mathrm{hi} \preccurlyeq x_u$, $f(\mathrm{m}) \succ y$, $f(\mathrm{hi}) \succ y$, $f\colon \mathbb{F} \to \mathbb{F}$, $y \in \mathbb{F}$, $[x_l, x_u] \in \mathcal{I}_\mathbb{F}$, $n_\mathrm{g} \geq n_\mathrm{g}^f$, $d_\mathrm{M} \geq d_\mathrm{M}^f$, $w_\mathrm{M} \geq w_\mathrm{M}^f$, $\alpha \preccurlyeq \alpha^f$, $\omega \succcurlyeq \omega^f$, $n_\mathrm{g} > 0$, $x_l \preccurlyeq \alpha \preccurlyeq \omega \preccurlyeq x_u$, $\alpha \leq \mathrm{hi} \wedge \omega \geq \mathrm{m}$, $t \in \mathbb{N}$.

Ensure: $b \in \{0, 1, 2\}$,
$n_\mathrm{g} = 1 \wedge w_\mathrm{M} < t \wedge \big(\alpha = \alpha^f \vee \omega = \omega^f \vee \big(\alpha \prec^k \omega \wedge k \leq t\big)\big) \implies b \in \{0, 1\}$,
$b = 0 \implies \forall x \in [\mathrm{m}, \mathrm{hi}] : f(x) \succ y$,
$b = 1 \implies z \in \mathbb{F} \wedge \mathrm{lo} \preccurlyeq z \preccurlyeq \mathrm{hi} \wedge \forall x \in (z, \mathrm{hi}] : f(x) \succ y \wedge f(z) \preccurlyeq y$.

1: **if** $n_\mathrm{g} = 1 \wedge w_\mathrm{M} \leq t$
$\wedge \big(f(\omega^-) \prec f(\omega) \vee f(\alpha^+) \prec f(\alpha) \vee \big(\alpha \prec^k \omega \wedge k \leq t\big)\big)$ **then**
2: $s_l := \max\{\alpha, \mathrm{lo}\}$
3: **if** $f(\omega^-) \prec f(\omega) \vee \big(\alpha \prec^k \omega \wedge k \leq t\big)$ **then**
4: $s_u := \min\{\omega, \mathrm{hi}\}$
5: **else**
6: $s_u := \min\{\alpha^{+w_\mathrm{M}}, \mathrm{hi}\}$
7: **end if**
8: $(b, z) :=$ `linsearch_leq`$(f, y, w_\mathrm{M}, s_l, s_u)$
9: $\triangleright\ (b = 0 \wedge z = \hat{x} \wedge \forall x \in [z, s_u] : f(x) \succ y)$
10: $\triangleright \vee (b = 1 \wedge z \in [\hat{x}, s_u] \wedge f(z) \preccurlyeq y \wedge \forall x \in (z, s_u] : f(x) \succ y)$
11: $\triangleright$ where $\hat{x} = \max\{s_l, s_u^{-w_\mathrm{M}}\}$
12: **else**
13: $b = 2$
14: **end if**

When $f(\mathrm{lo}) \succ y$, if the function has no glitches then we are sure its graph is completely below y, and equation $y = f(\mathbf{x})$ has no solution. Otherwise, there might be glitches in $[x_l, x_u]$ in which the function decreases until it reaches y. If there is only one glitch and it is not too wide, function `check_glitch` searches it float-by-float to find out whether the function actually reaches y. If not, we can state that our constraint is unsatisfiable; the appropriate predicate index is returned otherwise.

If lo and hi satisfy its precondition, Algorithm `bisect_ub` is eventually called on line 28. It has been designed to perform a dichotomic search in interval $[\mathrm{lo}, \mathrm{hi}]$, taking into account the fact that f may not be isotone due to glitches. In this function, `check_glitch` is used again in those cases where we do not know whether the function's graph is completely above y in the upper half of interval $[\mathrm{lo}, \mathrm{hi}]$ due to the possible presence of glitches. If performing the linear search is unfeasible because of the excessive number or width of glitches, function `logsearch_ub` tries to quickly find a value for hi such that the distance between $f(x)$ and y is too large for a glitch to be deep enough to let the function reach

Algorithm 7 Inverse propagation: `bisect_ub`$(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, s, t, \mathrm{lo}, \mathrm{hi})$

Require: $x_l \preccurlyeq \mathrm{lo} \prec \mathrm{hi} \preccurlyeq x_u$, $f(\mathrm{lo}) \preccurlyeq y \prec f(\mathrm{hi})$, $\forall x \in [\mathrm{hi}, x_u] : f(x) \succ y$ $f\colon \mathbb{F} \to \mathbb{F}$, $y \in \mathbb{F}$, $[x_l, x_u] \in \mathcal{I}_\mathbb{F}$, $n_\mathrm{g} \geq n_\mathrm{g}^f$, $d_\mathrm{M} \geq d_\mathrm{M}^f$, $w_\mathrm{M} \geq w_\mathrm{M}^f$, $\alpha \preccurlyeq \alpha^f$, $\omega \succcurlyeq \omega^f$, $n_\mathrm{g} > 0 \implies (x_l \preccurlyeq \alpha \preccurlyeq \omega \preccurlyeq x_u)$, $s, t \in \mathbb{N}$.

Ensure: Correctness: $x_l \preccurlyeq \mathrm{lo} \prec \mathrm{hi} \preccurlyeq x_u$, $f(\mathrm{lo}) \preccurlyeq y \prec f(\mathrm{hi})$, $\forall x \in [\mathrm{hi}, x_u] : f(x) \succ y$
Precision: $\big(n_\mathrm{g} = 0 \vee \big(n_\mathrm{g} = 1 \wedge w_\mathrm{M} < t \wedge \big(\alpha = \alpha^f \vee \omega = \omega^f \vee \big(\alpha \prec^k \omega \wedge k \leq t\big)\big)\big)\big) \implies f(\mathrm{hi}^-) \preccurlyeq y$

```
1: while hi ≻^1 lo do
2:     mid := split_point(lo, hi)    ▷ ∃m, m' > 0, |m − m'| ≤ 1, mid = hi^{−m} = lo^{+m'}
3:     if f(mid) ≼ y then
4:         lo := mid
5:     else if n_g = 0 ∨ hi ≼ α ∨ mid ≽ ω ∨ y ≺^{d_M} f(mid) then
6:         hi := mid
7:     else
8:         (b, z) := check_glitch(f, y, n_g, d_M, w_M, α, ω, t, lo, mid, hi)
9:         if b = 0 then
10:            hi := mid
11:        else if b = 1 then
12:            hi := z^+
13:            break
14:        else
15:            z := logsearch_ub(f, d_M, mid, hi, y, s)
16:                                   ▷ z ∈ [mid, hi] ∧ ((z ≺ hi) ⟹ f(z) ≻^{d_M} y)
17:            if z ≺ hi then
18:                hi := z
19:            else
20:                break
21:            end if
22:        end if
23:    end if
24: end while
```

y. Since this could make the partitioning uneven and undermine the logarithmic complexity of the process, the number of such calls of `logsearch_ub` is limited by threshold s.

Finally, on line 29 a **while** loop goes backwards float-by-float for a maximum of t steps, until it reaches a value for hi such that $f(\mathrm{hi}^-) \preccurlyeq y$, and returns it as the upper bound. If such value is not found, the algorithms returns the current value of hi with predicate p_7, that means an optimal value for the upper bound was not found.

The correctness of Algorithms 4, 5, 6 and 7 is formally proved in Section B.2 of the Appendix. The number of calls to function f performed at most by `upper_bound` has the form $k \log_2(\#[x_l, x_u]) + c$, where k and c are small constants that depend on w_M, s and t. The exact values of these constants in some special glitch data configurations are given in Theorems 5 and 6, in the Appendix.

5.3 Propagation Algorithms for Trigonometric Functions

The propagation algorithms for trigonometric functions exploit the fact that these functions are piece-wise quasi-isotonic or quasi-antitonic: the algorithms devised for regular quasi-monotonic functions are applied separately to each branch, in order to obtain refined intervals or multi-intervals. In order to split the domains of the functions according to the sets of the isotonic/antitonic intervals (see Definition 6), the techniques described in section 4.3 are used. In particular, interval $\mathbf{x} = [x_l, x_u]$ must be split in sub-intervals of the form $\left[\left[(k-2)\frac{\pi}{2}\right]_\uparrow, \left[k\frac{\pi}{2}\right]_\downarrow\right]$. The possible values of $k \in \mathbb{Z}$ depend on the function: we introduced $p \in \{\mathrm{even_v}, \mathrm{odd_v}, \mathrm{odd_c}\}$ to identify its behavior.

$\mathrm{even_v}$: the function changes its monotonicity in even multiples of $\frac{\pi}{2}$, i.e. in numbers of the type $2k\frac{\pi}{2}$, with $k \in \mathbb{Z}$. If $f = \cos$, this value of p is passed.

$\mathrm{odd_v}$: the function changes its tonicity in odd multiples of $\frac{\pi}{2}$, numbers of the form $(2k+1)\frac{\pi}{2}$, $k \in \mathbb{Z}$. This is the behavior of the sine function.

$\mathrm{odd_c}$: the function has a discontinuity in odd multiples of $\frac{\pi}{2}$, but it remains isotonic. The tangent has this behavior.

p can be used to identify points in which the monotonicity of the function changes. In order to identify the branch which a value $x \in \mathbb{F}$ belongs to, $k = \left\lceil x\frac{2}{\pi} \right\rceil$ is computed. Then, k is incremented if $p = \mathrm{even_v}$ and the initial value of k was odd, or if $p \in \{\mathrm{odd_v}, \mathrm{odd_c}\}$ and k was even. The obtained value of k is such that the function changes its monotonicity or has a discontinuity in $k\frac{\pi}{2}$ and $(k-2)\frac{\pi}{2}$, and $x \in \left[\left[(k-2)\frac{\pi}{2}\right]_\uparrow, \left[k\frac{\pi}{2}\right]_\downarrow\right]$.

Direct Propagation The periodicity of trigonometric functions poses a fundamental issue: in each monotonic branch the function can cover its whole range. Therefore, if the interval $[x_l, x_u]$ to be used to refine $\mathbf{y}$ spans multiple branches, almost no refinement can be done. However, since floating-point numbers become wider as the exponent grows, there is the possibility that some branches do not

reach the ends of the range, because there are no points where the function takes those values in such branches.

The algorithm that we devised for direct propagation in trigonometric functions tries to take advantage of these facts. First, it identifies the branches of the graph of the functions to which x_l and x_u belong, in the way described above. If they are both in the same monotonic branch, the refinement algorithm for regular functions described in Section 5.1 can be called. Otherwise, the said refinement function should be called separately for each branch, and then the minimum of the found values for y_l and the maximum of the found values for y_u should be returned. Since the number of branches to be separately inspected can be extremely high, a threshold g should be posed on the maximum number of branches to be analyzed. If $[x_l, x_u]$ spans more than g branches, the bounds of the function's range are returned.

Inverse Propagation The algorithm for inverse propagation, that we called `compute_bounds_trig`, refines the interval for `x` by splitting its initial domain $[x_l, x_u]$ into monotonic intervals, and then applying algorithms `lower_bound` and `upper_bound` locally in order to obtain a refined multi-interval. The multiplications by $\frac{\pi}{2}$ and $\frac{2}{\pi}$ needed by this algorithm are computed with sufficient precision, as explained in section 4.3.

Since interval $[x_l, x_u]$ can be wide, the number of sub-intervals it has to be split into can be excessively large. To overcome this issue, `compute_bounds_trig` accepts an argument g which defines a maximum number of sub-intervals to be returned. The intervals nearest to 0 are split with maximum granularity (each monotonic branch is considered separately), because floating-point numbers in this area are denser, and the probability of finding an optimal solution is higher. If necessary, the remaining branches are gathered in two larger intervals, one to the left of the domain and starting with x_l, and one to the right, ending with x_u. These intervals are only refined in branches at the boundaries. Obviously, if the domain does no cross 0, then only one of these two big sub-intervals is made, and the $g - 1$ branches nearest to 0 are treated individually in separate sub-intervals.

Functions `lower_bound` and `upper_bound` can be only called if $[x_l, x_u]$ is a subset of $[-\ell_{\max}, \ell_{\max}]$, because they cannot operate on sub-intervals whose length is too small (see section 4.1). If this algorithm is called repeatedly on the same interval, however, a value $g > 1$ could still cause complexity issues: if the single-branch intervals are discarded because no solutions are found in them, a repeated call of this refinement algorithm would split the multi-branch sub-interval again and again. To avoid this problem, the domain in which `compute_bounds_trig` is called can be further reduced.

`compute_bounds_trig` requires both y_l and y_u as arguments: y_l is used to refine the lower bounds, and y_u is used to refine the upper bounds. On quasi-antitonic intervals, function `lower_bound` is called with $-f$, which is quasi-isotonic in such intervals, and $f^{\mathrm{i}} \circ (-\mathrm{id})$ as the rough inverse of $-f$. In this cases, $-y_u$ is used instead of y_l. We proceed similarly with `upper_bound`.

The other arguments required by `compute_bounds_trig` are essentially the same as those for `lower_bound`, except for n_{gM}, which is a safe approximation of the maximum number of glitches in each quasi-monotonic branch of the function. For each branch, the value of n_{g} for the refinement function is set to 0 if the branch does not intersect $[\alpha, \omega]$, and to n_{gM} otherwise.

6 Implementation and Experiments

In this section we first describe the implementation of the algorithms described in this paper and its context, then, by means of an experiment, we show the kind of verification activities they make possible.

6.1 Implementation and Experiments

All the algorithms presented in this paper have been implemented and included in the ECLAIR software verification platform for C/C++ source code, Java source code and bytecode.[17] For the analysis of integers and floating-point values, ECLAIR uses different domains: mainly multi-intervals with a judicious use of polyhedral approximations made available by (the commercial version of) the Parma Polyhedra Library (PPL) [BHZ08]. For reasoning on the floating-point arithmetic operations, ECLAIR uses:

- algorithms realizing optimal direct projections as well as correct and precise indirect projections: the result is similar to the projections defined in [Mic02], but the ECLAIR algorithms never require working with precision greater than the operation data type;
- algorithms that exploit properties of the binary floating-point representations in order to obtain enhanced precision [BCGG16];
- dynamic linear relaxation techniques [BMR12,DGD07] using the PPL to enhance constraint propagation with the relational information provided by convex polyhedra.

The algorithms defined in this papers have been implemented in C++ and extensively tested on a variety of implementations with different characteristics in terms of the presence and nature of monotonicity glitches. These algorithms are now used in three components of the platform instantiation for C/C++: the semantic analysis engine based on abstract interpretation [CC77], the automatic generator of test inputs, and the symbolic model checker, the latter being both based on constraint solving [GBR98,GBR00]. All components use multi-interval refinement, though in different ways: the test generator and symbolic model checker are driven by labeling and backtracking search. As these have a negative interaction with the searches controlled by the s and t parameters of the algorithms, their setting needs to be controlled more carefully (and they are better set to 0 when glitch data is precise) by these components, whereas they are used

[17] `http://bugseng.com/products/eclair`

with values in the range 5–20 in the semantic analysis engine. Here, we only report about experimentation with the symbolic model checker and automatic generator of test input. One interesting feature of the latter is that it optionally produces a transformed source program that contains the original program, suitably instrumented, and a driver that runs one of the generated tests at a time. The instrumented code checks that each one of the generated tests achieves its target, e.g., it reaches a certain program point, or it causes an integer overflow or the generation of a floating-point infinite value. The validation of test inputs is thus completely automatic.

6.2 Experiments

The development of the algorithms of Section 5 has been paired by the development of a test suite to exercise them in the most disparate ways. However, to better illustrate the potentiality such algorithms have, let us consider again our introductory example, which is reproduced in Listing 1.2, suitably annotated. Let us pretend we know nothing about such code (which is realistic, as there are no comments besides the one reported). So, we initially assume that the entry point is `latlong_utm_of()`; as there are no assertions, we also assume all inputs are possible. For an exploratory analysis, we use ECLAIR's symbolic model checker in order to detect the possible presence of run-time anomalies: overflow, division be zero and other source of undefined and implementation-defined behavior over the integers, inexact integer-to-floating conversions, finite-to-infinite and numeric-to-NaN transitions over floating point numbers. A *finite-to-infinite* (resp., *numeric-to-NaN*) transition is a computation whereby the inputs to a floating-point operation or `math.h`/`cmath` function is finite (resp., numeric) and the output is infinite (resp, NaN). We also set an analysis parameter asking ECLAIR to flag all the invocations of trigonometric functions whose argument has an absolute value greater than, say, 16. Not surprisingly, we obtain three test inputs showing that this is indeed possible; they concern the following program points:

p_2 : $(-\texttt{0x86487f.p} - \texttt{18F}, +\texttt{0.0F}, \texttt{1})$, where $-\texttt{0x86487f.p} - \texttt{18F} \approx -33.570797$,
p_3 : $(-\texttt{0x8a3ae7.p} - \texttt{19F}, +\texttt{0.0F}, \texttt{1})$, where $-\texttt{0x8a3ae7.p} - \texttt{19F} \approx -17.278761$,
p_5 : $(+\texttt{0.0F}, -\texttt{0x98b6c1.p} - \texttt{19F}, \texttt{1})$, where $-\texttt{0x98b6c1.p} - \texttt{19F} \approx -19.089235$.

Of course, the latter input causes the same phenomenon at program point p_7 as well. Fair enough: perhaps `latlong_utm_of()` callers only pass smaller values for `phi` and `lambda`; even if that is not the case, then perhaps the only problem is a slight precision issue. But ECLAIR produces two other test inputs, with the specification that they trigger number-to-NaN transitions:

p_1 : $(+\texttt{0xc90fdb.p} - \texttt{23F}, +\texttt{0.0F}, \texttt{1})$, where $+\texttt{0xc90fdb.p} - \texttt{23F} \approx 1.570796$,
p_4 : $(-\texttt{0x8a9483.p} - \texttt{53F}, +\texttt{0xcfc398.p} - \texttt{23F}, \texttt{1})$, where $-\texttt{0x8a9483.p} - \texttt{53F} \approx -1.0083 \cdot 10^{-9}$, and $+\texttt{0xcfc398.p} - \texttt{23F} \approx -17.278761$.

As $+\texttt{0xc90fdb.p} - \texttt{23F} \approx 1.570796$ converted to double precision is slightly greater than `M_PI_2`, the round-to-nearest, double-precision approximation of

Listing 1.2. Code excerpted from a real-world avionic library, annotated

```
#include <math.h>
#include <stdint.h>

#define RadOfDeg(x) ((x) * (M_PI/180.))
#define E 0.08181919106 /* Computation for the WGS84 geoid only */
#define LambdaOfUtmZone(utm_zone) RadOfDeg((utm_zone-1)*6-180+3)
#define CScal(k, z) { z.re *= k; z.im *= k; }
#define CAdd(z1, z2) { z2.re += z1.re; z2.im += z1.im; }
#define CSub(z1, z2) { z2.re -= z1.re; z2.im -= z1.im; }
#define CI(z) { float tmp = z.re; z.re = - z.im; z.im = tmp; }
#define CExp(z) { float e = exp(z.re); z.re = e*cos(z.im); \
  z.im = e*sin(z.im); }
#define CSin(z) { CI(z); struct complex _z = {-z.re, -z.im}; \
  float e = exp(z.re); float cos_z_im = cos(z.im); z.re = e*cos_z_im; \
  float sin_z_im = sin(z.im); z.im = e*sin_z_im; _z.re = cos_z_im/e; \
  _z.im = -sin_z_im/e; CSub(_z, z); CScal(-0.5, z); CI(z); }

static inline float isometric_latitude(float phi, float e) {
  return log^p1(tan^p2(M_PI_4 + phi / 2.0))
    - e / 2.0 * log((1.0 + e * sin^p3(phi)) / (1.0 - e * sin(phi)));
}

static inline float isometric_latitude0(float phi) {
  return log^p4(tan(M_PI_4 + phi / 2.0));
}

void latlong_utm_of(float phi, float lambda, uint8_t utm_zone) {
  float lambda_c = LambdaOfUtmZone(utm_zone);
  float ll = isometric_latitude(phi, E);
  float dl = lambda - lambda_c;
  float phi_ = asin(sin^p5(dl) / cosh(ll));
  float ll_ = isometric_latitude0(phi_);
  float lambda_ = atan(sinh(ll) /^p6 cos^p7(dl));
  struct complex z_ = { lambda_, ll_ };
  CScal(serie_coeff_proj_mercator[0], z_);
  uint8_t k;
  for(k = 1; k < 3; k++) {
    struct complex z = { lambda_, ll_ };
    CScal(2*k, z);
    CSin(z);
    CScal(serie_coeff_proj_mercator[k], z);
    CAdd(z, z_);
  }
  CScal(N, z_);
  latlong_utm_x = XS + z_.im;
  latlong_utm_y = z_.re;
}
```

Listing 1.3. Calling context of `latlong_utm_of()`

```
1  int gps_lat, gps_lon; /* 1e7 deg */
2  unsigned char nav_utm_zone0;
3  /* [...] */
4  static gboolean
5  read_data(GIOChannel *chan, GIOCondition cond, gpointer data) {
6    int count;
7    char buf[BUFSIZE];
8
9    /* receive data packet containing formatted data */
10   count = recv(sock, buf, sizeof(buf), 0);
11   if (count > 0) {
12     if (count == 23) {
13 // FillBufWith32bit(com_trans.buf, 1, gps_lat);
14       gps_lat = buf2uint(&buf[0]);
15 // FillBufWith32bit(com_trans.buf, 5, gps_lon);
16       gps_lon = buf2uint(&buf[4]);
17       /* [...] */
18
19
20       nav_utm_zone0 =* (gps_lon/10000000+180) / 6 + 1;
21       latlong_utm_of(RadOfDeg(gps_lat/1e7), RadOfDeg(gps_lon/1e7),
22                      nav_utm_zone0);
```

$\frac{\pi}{2}$ defined in `math.h`, we make the hypothesis that `phi` has to be less than or equal to `M_PI_2`. Indeed, looking at the function callers (there is only one in the program), we come to the realization that `phi` and `lambda` are a latitude and longitude in radians, respectively. We attempt validation of this hypothesis by adding the assertions

```
assert(-M_PI_2 <= phi && phi <= M_PI_2);
assert(-M_PI <= lambda && lambda <= M_PI);
```

at the beginning of `latlong_utm_of()` and repeat the analysis. We obtain another ill-conditioned trigonometric function argument test input for program point p_5:

p_5 : $(+0.0\texttt{F}, +0\texttt{xc90fda.p} - 22\texttt{F}, 255)$, where $-0\texttt{x98b6c1.p} - 19\texttt{F} \approx 3.141593$

(surely `utm_zone` $= 255$ is not among the expected inputs) and another numeric-to-NaN transition input:

p_4 : $(-0\texttt{x8a1e4e.p} - 53\texttt{F}, +0\texttt{xc5b5fe.p} - 22\texttt{F}, 15)$, where $-0\texttt{x8a1e4e.p} - 53\texttt{F} \approx -1.00494 \cdot 10^{-9}$, and $+0\texttt{xc5b5fe.p} - 22\texttt{F} \approx 3.089233$.

In order to understand what are the intended inputs for `latlong_utm_of()` we now take into account its calling context, summarized in Listing 1.3. Basically, the inputs to `latlong_utm_of()` depend on two 32-bit signed integers, `gps_lat`

and `gps_lon`, that are received from a communication channel: no check is made upon them after reading the values out of the input buffer. Taking into account the caller context, ECLAIR generates three reports: if `gps_lat` $= 0$ and `gps_lon` $= -1920000000$ at line 20 in Listing 1.3, then the conversion in the assignment marked with '$*$' on the same line causes an unsigned wraparound ($-1 \bmod 256 = 255$, so that, yes, `latlong_utm_of()` can be called with `utm_zone` $= 255$); the same input also generates an ill-conditioned trigonometric function argument for program point p_5 in Listing 1.2; most importantly, if `gps_lon` $= 900000059$ and `gps_lat` $= -1920000000$, then we have a numeric-to-NaN transition at program point p_1. This probably means that if the equipement at the other hand of the communication channel is defective or if there is a communication error, things can go horribly wrong. However, let us now suppose that there are no problems of this kind and that we have $|\texttt{gps_lat} \cdot 10^{-7}| \leq 90$ and $|\texttt{gps_lon} \cdot 10^{-7} \leq 180$ as the code seems to assume. The analysis with ECLAIR shows this is not enough: the numeric-to-NaN transition at program point p_1 is still possible with `gps_lat` $= -899999991$ and `gps_lon` $= -1800000000$.[18] In a couple more iterations we add the assertions

```
18'   assert( -899999990 <= gps_lat && gps_lat <=  899999990);
19'   assert(-1800000000 <= gps_lon && gps_lon <= 1800000000);
```

and the final ECLAIR run shows no report. This, per se, does not mean much. However, this experiment was done on the `zoltan` machine, for which we have precise glitch data for the single-precision functions (which are not used in the code considered) and we have the maximum known errors provided by the GNU libc manual for the double-precision functions [LSM+16]. As explained in Section 7.2, this data provides imprecise and possibly incorrect information about glitches that our algorithms can exploit. In turn, all this means that:

- *if* the numbers in [LSM+16] do really provide upper bounds to the maximum errors of the used functions, and
- *if* the caller guarantees that the values of `gps_lat` and `gps_lon` do satisfy the "stay away from the poles" assertions at lines 18′–19′ in modified Listing 1.3,
- *then*, in the context of such call, all the 154 potential run-time anomalies[19] in the 90 potentially problematic program points of Listing 1.2 cannot occur on `zoltan`.

Just to mention one potential problem, division by zero and consequent finite-to-infinite transition at program point p_6, cannot happen on that implementation.

7 Discussion and Further Work

In this Section we discuss some aspects related to the applicability of our proposal: every such aspect immediately suggests directions for further work.

[18] This point is roughly 10 cm from the Geographic South Pole.

[19] More precisely: 4 integer overflows, 4 inexact conversions, 10 ill-conditioned trigonometric function arguments, 70 finite-to-infinity and 66-numeric-to-NaN transitions.

7.1 Access to the Target Library

For the purposes of true verification, our approach requires execution access to the mathematical library used by the target. Of course, when the host and the target computer can coincide, i.e., when the target is powerful enough to run the verifier code, this is no problem. Alternatively, the host computer might provide an implementation that is fully equivalent to the one used on the target computer: this is the case, e.g., of targets where all floating-point support is implemented in software. In other cases, using an emulator is required. Luckily, emulation technology is improving very rapidly due to the increasing adoption of virtualization.

This can be seen as the major drawback of our approach. However, in our opinion the question should be put in the following terms: in order to properly verify a piece of code using library functions against, say, the abstence of run-time anomalies, the library functions have to be fully specified. If a specification of the form "all functions are POSIX-compliant [IEE13] and compute correctly-rounded results" is available, then fine. Otherwise there really is no other way than supplementing the partial specification available with the missing bits: providing execution access to the library during the analysis along with correct bounds on the size of the glitches might well be the less expensive option.

7.2 Obtaining Glitch Data

The other requirement of the approach described in this paper concerns the availability of (possibly imprecise) information about the monotonicity glitches. Some ways in which such information can (at least in principle) be obtained are the topics of the next sections.

Brute Force For single-precision IEEE 754 (unary) functions collection of precise glitch data by brute force is perfectly feasible. Glitch data presented in this paper have been obtained in this way. For the 25 functions studied in this paper, it takes less than two hours on ordinary PC hardware. With current less powerful CPUs used on embedded systems, it might take ten or twenty times as much. This is not really a problem as glitch data must be collected only once for each implementation of the `math.h`/`cmath` functions. And, especially in safety-critical sectors, the mathematical (and other) libraries will rarely if ever be changed once they have been selected. Of course, this method cannot be used for double-precision or extended-precision functions.

Using Precision Guarantees When Available When the target mathematical library comes equipped with information on the maximum errors for each function, such information can be used to determine safe approximations of the glitch parameters required by our algorithms. Given a function and an architecture, the maximum error is measured in ULP and can be used as an upper bound for the maximal depth of the glitches w_M. Given an interval $[x_l, x_u]$, in

absence of more precise information, the cardinality of the floating-point interval $[x_l, x_u]$ is as an upper bound to the maximum width of the glitches and, of course x_l and x_u are safe approximations of where the glitches begin and end, respectively. Finally, setting $n_g > 1$ allows us to call the indirect propagation algorithms to refine the interval $[x_l, x_u]$ of the function domain with respect to a given interval $[y_l, y_u]$ of the function range. Even with such a rough information on the glitches, the algorithms would allow us to refine the interval $[x_l, x_u]$ using the logarithmic searches (`logsearch_lb` and `logsearch_ub`) and returning $[x'_l, x'_u] \subseteq [x_l, x_u]$ such that $f(x'_l)$ is smaller than y_l more than w_M ULPs and $f(x'_u)$ is bigger than y_u more than w_M ULPs. Therefore, the cardinality of the resulting refined interval is strictly related to the speed of growth of the considered function. For future work, we intend to investigate how information on the maximum error, coupled with the knowledge of the function and of the interval to be refined, allow computing sound and tight bounds to the width of glitches on that interval.

Analysis of the Implementation The transcendental functions are usually implemented computing polynomial approximations. When speed is more important than precision, such computations are carried out in the same floating-point format as the the function being approximated; otherwise extended precision can be used to reduce the error. Whereas the total error accumulation can be bounded, the ordinary techniques used to do so do not allow relating the rounding errors for different input values to one another. So, if the error bound is small enough to imply monotonicity, fine. Otherwise, as things stand today, we are left with the approach of the previous section. However, we conjecture that (some of) the implementation algorithms can be analyzed with other techniques in order to obtain mode precise glitch data: this is another direction for future work.

7.3 Remaining Functions in `math.h/cmath`

In this paper we described how to deal with 75 of the standard C/C++ mathematical functions: but there are many others. Several of them are not problematic, as they are fully specified and their treatment poses no problem (e.g., the families of operations `round`, `trunc`, `floor`, `ceil`, `fma`, `fabs`, `floor`, `ceil`, `next`, ...). In future work we will focus on the remaining functions, particularly the functions with two inputs, such as `atan2` and `pow`, and then the complex functions.

7.4 The Case of Correctly Rounded Libraries

When a function is correctly rounded for the rounding modes of interest, then the approach presented in this paper is clearly applicable: all the quasi-monotonic branches are also monotonic, zero glitches. However, if a library provides correctly rounded versions of both $f\colon \mathbb{R} \to \mathbb{R}$ and $f^{-1}\colon \mathbb{R} \to \mathbb{R}$ for all rounding modes, monotonicity allows a much more efficient approach: [Mic02] presents

optimal projections for isotonic unary functions over one argument that can be generalized to all piecewise monotonic functions.

7.5 Better Labeling Strategies for Constraint-Based Reasoning

Constraint-based symbolic model checkers and test input generators are based on search algorithms that operate by interleaving two rather different processes:

1. *constraint propagation*, in which constraints are used to refine the domains of variables (intervals or multi-intervals in our case), and
2. *labeling*, whereby a variable is chosen and its domain is partitioned into two or more subsets, each of which is explored separately.

It is the second process that drives the first one: when constraint propagation goes to *quiescence*, i.e., when no further refinement of the domains can be achieved, labeling splits the domain of the chosen variable, and this triggers a new phase of constraint propagation. This goes on until a solution has been found or one of the domains becomes empty, in which case unsatisfiability of the original problem is proved.

In this paper we only dealt with constraint propagation, but different choices, as far as the labeling strategies are concerned, have an enormous influence on performance. Unfortunately, there is no such a thing as *the* good labeling strategy (though it is possible to recognize some definitely bad labeling strategies): it is a matter of heuristics, and strategies that work well for one problem may still work badly for another problem. It is worth observing that test input generation and symbolic model checking give rise to constraint problems that may be of very different nature: while the latter are very often over-constrained (that is, there are few or no solutions at all, as the program in question exhibits very few or no run-time anomalies), this is not the case for the former (e.g., in the extreme case, a function body that is a basic block can be covered by a single test-input chosen more or less at random). This implies that input generation and symbolic model checking can both profit from the choice of different labeling strategies.

During the experimental part of our work on mathematical functions, we got a strong feeling that the current labeling strategy employed by ECLAIR can be significantly enhanced by defining heuristics that take into account how variables are constrained by invocations to such functions. Work on these new heuristics is ongoing.

8 Conclusion

There is a popular quotation in the software verification and validation community, whereby "Without a specification, a system cannot be right or wrong, it can only be surprising!"[20] This captures quite well the current state of affairs for C/C++ software that uses the functions declared in the standard `math.h`/`cmath`

[20] Paraphrased from [YBK85].

header files. Despite the progress made on the development of correctly rounded functions,[21] all implementations in widespread use, especially in the world of embedded systems, offer little or no guarantees about the computed results. As a consequence, the verification of programs using such functions is always painful and expensive and, for these reasons, more often than not it is only partially performed through testing. As the search space can be huge, testing can only cover a tiny fraction of all the possible value combinations: this cannot exclude the manifestation of unexpected results, certainly not with the level of confidence that is required for mission- and safety-critical applications.

The aim of this work is to improve upon the current situation now, i.e., without waiting for the wider adoption of correctly-rounded implementations. While such adoption is generally desirable and will certainly take place, at some stage and in some application domains, it is not clear whether correctly-rounded implementations can meet the efficiency criteria of all application domains, particularly in the field of embedded systems. Studying different implementations of the standard C/C++ mathematical functions, we realized that what they have in common is a piecewise quasi-monotonicity property: monotonicity is either preserved or only perturbed by small and, on average, not too frequent "glitches". Based on this observation, we developed direct and indirect propagation algorithms for the refinement of intervals and multi-intervals. These algorithms can be integrated into abstract interpreters as well as symbolic model checkers and automatic test input generators based on constraint propagation.

The techniques proposed here are now used in the C/C++ semantic analysis components of the ECLAIR software verification platform and the initial experiences are quite positive. We are now able to properly verify absence of run-time anomalies for code using the C/C++ standard functions that, before, was completely out of reach. Verification in the strong sense is only feasible modulo the possibility of bounding the size of glitches (this can always be done for the single-precision functions) and the possibility to query the underlying implementation of the functions during the analysis. For the cases where the first condition cannot be granted, we can still detect many definite program issues, even though we cannot draw conclusions from the fact issues have not been found. When the second condition cannot be met, still it may be possible to use a reference implementation with significant commonalities with the target implementation (the case where libraries for different architectures are derived from the same code base is quite common), and we can nonetheless detect high-severity, possible program issues.

Even though much remains to be done before we can claim that the problem of the verification of C/C++ programs using the standard mathematical functions has been solved from the practical point of view, we believe the present work is a definite step forward in the right direction, and one that has the potential of improving, starting from today, the current state of the art.

[21] See, e.g., the very interesting MetaLibm project at `http://www.metalibm.org/` (last accessed on October 23rd, 2016).

Acknowledgments We are grateful to Arnaud Gotlieb (Simula Research Laboratory, Norway, and INRIA – Rennes, France) and Claude Michel (INRIA – Sophia Antipolis, France) for the fruitful discussions we had on the subject of this paper. We are are also grateful to Marcel Beemster (Solid Sands, The Netherlands) for the discussions we had on the subject of testing mathematical libraries for the C language. Alessandro Zaccagnini (University of Parma, Italy) helped us defining rough inverses for the Gamma functions. Thanks also to Patricia M. Hill (BUGSENG srl, Italy) for her work on the ECLAIR platform that greatly facilitated the experimental evaluation of this work.

References

BCGG16. R. Bagnara, M. Carlier, R. Gori, and A. Gotlieb, *Exploiting binary floating-point representations for constraint propagation*, INFORMS Journal on Computing **28** (2016), no. 1, 31–46.

BDL12. L. Burdy, J.-L. Dufour, and T. Lecomte, *The B method takes up floating-point numbers*, Proceedings of the 6th International Conference & Exhibition on Embedded Real Time Software and Systems (ERTS 2012) (Toulouse, France), 2012, Available at `http://web1.see.asso.fr/erts2012/Site/0P2RUC89/5C-2.pdf`.

BHZ08. R. Bagnara, P. M. Hill, and E. Zaffanella, *The Parma Polyhedra Library: Toward a complete set of numerical abstractions for the analysis and verification of hardware and software systems*, Science of Computer Programming **72** (2008), no. 1–2, 3–21.

BKD+05. N. Brisebarre, P. Kornerup, D. Defour, N. Revol, and J.-M. Muller, *A new range-reduction algorithm*, IEEE Transactions on Computers **54** (2005), 331–339.

BMR12. M. S. Belaid, C. Michel, and M. Rueher, *Boosting local consistency algorithms over floating-point numbers*, Proceedings of the 18th International Conference on Principles and Practice of Constraint Programming (Québec City, Canada) (M. Milano, ed.), Lecture Notes in Computer Science, vol. 7514, Springer-Verlag, Berlin, 2012, pp. 127–140.

CC77. P. Cousot and R. Cousot, *Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints*, Proceedings of the Fourth Annual ACM Symposium on Principles of Programming Languages (Los Angeles, CA, USA), ACM Press, 1977, pp. 238–252.

Dek71. T. J. Dekker, *A floating-point technique for extending the available precision*, Numerische Mathematik **18** (1971), no. 3, 224–242.

DGD07. T. Denmat, A. Gotlieb, and M. Ducassé, *Improving constraint-based testing with dynamic linear relaxations*, Proceedings of the 18th IEEE International Symposium on Software Reliability (ISSRE 2007) (Trollhättan, Sweden), IEEE Computer Society, 2007, pp. 181–190.

DMMM96. M. Daumas, C. Mazenc, X. Merrheim, and J.-M. Muller, *Modular range reduction: a new algorithm for fast and accurate computation of the elementary functions*, J.UCS The Journal of Universal Computer Science: Annual Print and CD-ROM Archive Edition Volume 1 • 1995 (Hermann Maurer, Cristian Calude, and Arto Salomaa, eds.), Springer Berlin Heidelberg, Berlin, Heidelberg, 1996, pp. 162–175.

GBR98. A. Gotlieb, B. Botella, and M. Rueher, *Automatic test data generation using constraint solving techniques*, Proceedings of the 1998 ACM SIGSOFT International Symposium on Software Testing and Analysis (New York, NY, USA), ISSTA '98, ACM, 1998, pp. 53–62.

GBR00. ———, *A CLP framework for computing structural test data*, Computational Logic — CL 2000: First International Conference London, UK, July 24–28, 2000 Proceedings, Springer Berlin Heidelberg, 2000, pp. 399–413.

HLB01. Y. Hida, X. S. Li, and D. H. Bailey, *Algorithms for quad-double precision floating point arithmetic*, Proceedings of the 15th IEEE Symposium on Computer Arithmetic, IEEE Computer Society, August 2001, pp. 155–162.

IEE08. IEEE Computer Society, *IEEE standard for floating-point arithmetic*, IEEE Std 754-2008 (revision of IEEE Std 754-1985) ed., August 2008.

IEE13. IEEE Computer Society and The Open Group, New York, NY, USA, *Standard for Information Technology— Portable Operating System Interface (POSIX®), Base Specifications, Issue 7*, IEEE Std 1003.1, 2013 (incorporates IEEE Std 1003.1-2008, and IEEE Std 1003.1-2008/Cor 1-2013) ed., April 2013.

Int11. International Organization for Standardization, *ISO/IEC 9899:2011: Programming languages — C*, International Organization for Standardization, Geneva, Switzerland, 2011.

Int14. ———, *ISO/IEC 14882:2014(E): Programming languages — C++*, International Organization for Standardization, Geneva, Switzerland, 2014.

KM84. W. Kahan and S. McDonald, *Minimizing* `q * m - n`, unpublished handout (available at `https://people.eecs.berkeley.edu/~wkahan/testpi`), 1984.

LMT98. V. Lefèvre, J.-M. Muller, and A. Tisserand, *The table maker's dilemma*, Rapport de recherche 98–12, Laboratoire de l'Informatique du Parallélisme, École Normale Superiéure de Lyon, 1998.

LSM+16. S. Loosemore, R. M. Stallman, R. McGrath, A. Oram, and U. Drepper, *The GNU C library reference manual*, Free Software Foundation, Inc., Boston, MA, USA, version 2.23 ed., August 2016.

Mic02. C. Michel, *Exact projection functions for floating point number constraints*, Proceedings of the 7th International Symposium on Artificial Intelligence and Mathematics (Fort Lauderdale, FL, USA), 2002.

MPT16. J. M. Muller, V. Popescu, and P. T. P. Tang, *A new multiplication algorithm for extended precision using floating-point expansions*, 2016 IEEE 23nd Symposium on Computer Arithmetic (ARITH), IEEE Computer Society, July 2016, pp. 39–46.

Mul05. J.-M. Muller, *On the definition of* ulp(x), Rapport de recherche 5504, INRIA, 2005.

Ng92. K. C. Ng, *Argument reduction for huge arguments: Good to the last bit*, Tech. report, SunPro – A Sun Microsystems, Inc. Business, July 1992.

PH83. M. H. Payne and R. N. Hanek, *Radian reduction for trigonometric functions*, SIGNUM Newsletter **18** (1983), no. 1, 19–24.

TDB+06. S. Tucker Taft, R. A. Duff, R. L. Brukardt, E. Plödereder, and P. Leroy (eds.), *Ada 2005 Reference Manual. Language and Standard Libraries: International Standard ISO/IEC 8652/1995 (E) with Technical Corrigendum 1 and Amendment 1*, Lecture Notes in Computer Science, vol. 4348, Springer-Verlag, Berlin, Germany / Heidelberg, Germany / London, UK / etc., 2006.

VSD98. P. Van Hentenryck, V. Saraswat, and Y. Deville, *Design, implementation, and evaluation of the constraint language* `cc(fd)`, The Journal of Logic Programming **37** (1998), no. 13, 139 – 164.

YBK85. W. D. Young, W. E. Boebert, and R. Y. Kain, *Proving a computer system secure*, The Scientific Honeyweller **6** (1985), no. 2, 18–27.

A Glitch Data for Other Implementations of `libm`

Table 4. Glitch data for the `alpha` machine

function	$D_{\min}$	D_{M}	near			up			down			zero		
			n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}
`acosf`	-1	1												
`acoshf`	1	∞							1	1	2	1	1	2
`asinf`	-1	1												
`asinhf`	$-\infty$	∞							2	1	2	2	1	2
`atanf`	$-\infty$	∞				1	1	10^8						
`atanhf`	-1	1							2	1	2	2	1	2
`cbrtf`	$-\infty$	∞	10^6	1	2	10^6	1	2	10^6	1	2	10^6	1	2
`coshf`	$-\infty$	∞	454	1	2	466	1	2	442	1	2	448	1	2
`erff`	$-\infty$	∞												
`expf`	$-\infty$	∞												
`exp10f`	$-\infty$	∞												
`exp2f`	$-\infty$	∞	1	1	2	3	1	2	2	1	2	1	1	2
`expm1f`	$-\infty$	∞												
`lgammaf`	2	∞	163	1	2	164	1	2	166	1	2	161	1	2
`logf`	0	∞												
`log10f`	0	∞												
`log1pf`	-1	∞							1	1	2	1	1	2
`log2f`	0	∞												
`sinhf`	$-\infty$	∞												
`sqrtf`	0	∞												
`tanhf`	$-\infty$	∞				1	1	2	2	1	3			
`tgammaf`	2	∞	10^4	2	3	10^4	2	4	10^4	3	3	10^4	3	4
`cosf`	-2^{23}	2^{23}												
`sinf`	-2^{23}	2^{23}												
`tanf`	-2^{23}	2^{23}												

Table 5. Glitch data for the `gcc111` machine

function	$D_{\min}$	D_{M}	near			up			down			zero		
			n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}
`acoshf`	1	∞												
`asinhf`	$-\infty$	∞												
`atanf`	$-\infty$	∞												
`atanhf`	-1	1												
`coshf`	$-\infty$	∞												
`erff`	$-\infty$	∞												
`lgammaf`	2	∞												
`sinhf`	$-\infty$	∞												
`tanhf`	$-\infty$	∞				2	1	10^7						
`tgammaf`	2	∞												
`cosf`	-2^{23}	2^{23}												
`sinf`	-2^{23}	2^{23}												
`tanf`	-2^{23}	2^{23}												

Table 6. Glitch data for the `igor` machine

function	$D_{\min}$	D_{M}	near			up			down			zero		
			n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}
`acosf`	-1	1							1	1	10^{10}	1	1	10^{10}
`acoshf`	1	∞							1	1	2	1	1	2
`asinf`	-1	1												
`asinhf`	$-\infty$	∞							2	1	2	2	1	2
`atanf`	$-\infty$	∞												
`atanhf`	-1	1							2	1	2	2	1	2
`cbrtf`	$-\infty$	∞	10^6	1	2	10^6	1	2	10^6	1	2	10^6	1	2
`coshf`	$-\infty$	∞	454	1	2	466	1	2	442	1	2	448	1	2
`erff`	$-\infty$	∞												
`expf`	$-\infty$	∞												
`exp10f`	$-\infty$	∞												
`exp2f`	$-\infty$	∞	1	1	2	3	1	2	2	1	2	1	1	2
`expm1f`	$-\infty$	∞												
`lgammaf`	2	∞	163	1	2	164	1	2	166	1	2	161	1	2
`logf`	0	∞												
`log10f`	0	∞												
`log1pf`	-1	∞							1	1	2	1	1	2
`log2f`	0	∞												
`sinhf`	$-\infty$	∞												
`sqrtf`	0	∞												
`tanhf`	$-\infty$	∞				1	1	2	2	1	3			
`tgammaf`	2	∞	155	109	2	155	122	2	157	119	2	153	119	2
`cosf`	-2^{23}	2^{23}	10^4	1	3	10^4	1	3	10^4	1	3	10^4	1	3
`sinf`	-2^{23}	2^{23}												
`tanf`	-2^{23}	2^{23}												

Table 7. Glitch data for the `raspi` machine

function	$D_{\min}$	D_{M}	near			up			down			zero		
			n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}	n_{g}	d_{M}	w_{M}
`acosf`	-1	1												
`acoshf`	1	∞							1	1	2	1	1	2
`asinf`	-1	1												
`asinhf`	$-\infty$	∞							2	1	2	2	1	2
`atanf`	$-\infty$	∞				1	1	10^8						
`atanhf`	-1	1							2	1	2	2	1	2
`cbrtf`	$-\infty$	∞	10^6	1	2	10^6	1	2	10^6	1	2	10^6	1	2
`coshf`	$-\infty$	∞	454	1	2	466	1	2	442	1	2	448	1	2
`erff`	$-\infty$	∞												
`expf`	$-\infty$	∞												
`exp10f`	$-\infty$	∞												
`exp2f`	$-\infty$	∞	1	1	2									
`expm1f`	$-\infty$	∞												
`lgammaf`	2	∞	163	1	2	164	1	2	166	1	2	161	1	2
`logf`	0	∞												
`log10f`	0	∞												
`log1pf`	-1	∞							1	1	2	1	1	2
`log2f`	0	∞												
`sinhf`	$-\infty$	∞												
`sqrtf`	0	∞												
`tanhf`	$-\infty$	∞				1	1	2	2	1	3			
`tgammaf`	2	∞	10^4	2	3	10^4	2	4	10^4	3	3	10^4	3	4
`cosf`	-2^{23}	2^{23}	10^4	1	3	10^4	1	3	10^4	1	3	10^4	1	3
`sinf`	-2^{23}	2^{23}												
`tanf`	-2^{23}	2^{23}												

B Proofs

B.1 Computation of Lower Bounds

Lemma 2. *Function* `findhi_lb` *specified in* Algorithm 2 *satisfies its contract.*

Proof. We assume the precondition for `findhi_lb`$(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, t)$ is satisfied and, in particular $f(x_u) \prec y$. We have the following cases in which the condition on line 2 is true:

$n_\mathrm{g} = 0$: as $n_\mathrm{g} \geq n_\mathrm{g}^f$ we have $n_\mathrm{g}^f = 0$, that is, f is isotonic on $[x_l, x_u]$; thus $f(x_u) \prec y$ implies $f(x) \prec y$ for each $x \in [x_l, x_u]$.

$x_u \succ \omega$: as $\omega \geq \omega^f$, we know that $x_u \succ \omega^f$; this implies that x_u cannot be inside a glitch, since, by definition of ω^f, f is isotonic on $[\omega^f, x_u]$; therefore $f(x) \preccurlyeq f(x_u) \prec y$ for each $x \in [x_l, x_u]$.

$y \succ^{d_\mathrm{M}} f(x_u)$: this means that $f(x_u)$ is so low that, even under the worst-case assumption x_u is the minimum of a maximal-depth glitch, the value of f just before the glitch would still be below y; therefore, $f(x) \prec y$ for each $x \in [x_l, x_u]$.

In all circumstances in which the condition on line 2 is true, the post-condition for the case $r = 0$ is proved.

At line 4, since we are on the else statement, we know that $n_\mathrm{g} > 0$, $x_u \preccurlyeq \omega$ and $y \preccurlyeq^{d_\mathrm{M}} f(x_u)$. Assume further that the condition on line 4 is true, hence $n_\mathrm{g} = 1$. There are two cases:

$y \succ f(\alpha)$: this case is further split into two cases:

- $f(\alpha^+) \prec f(\alpha)$: in this case, f has exactly one glitch in $[x_l, x_u]$ beginning at $\alpha^f = \alpha$; hence $f(x) \preccurlyeq f(\alpha^f)$ for each $x \in [x_l, \alpha^f]$ and, since $y \succ f(\alpha)$, we also have $f(x) \prec y$ for each $x \in [x_l, \alpha^f]$; moreover, the algorithm precondition entails that $f(x_u) \prec y$, which, together with $n_\mathrm{g}^f = 1$, allows us to conclude that $f(x) \prec y$ also in the interval $(\alpha^f, x_u]$; summing up, $[x_l, x_u]$ does not contain any solutions for $y = f(\mathbf{x})$ and setting $r = 0$ in line 7 satisfies the post-condition.
- $f(\alpha^+) \succcurlyeq f(\alpha)$: in this case, f either has no glitch or it has exactly one glitch strictly past α (that is, $\alpha^f \succ \alpha$); as glitches can be too wide (i.e., $w_\mathrm{M} > t$), we refrain from searching a suitable value for hi to start bisection; as $y \succ f(\alpha)$, we have $y \succ f(x)$ for each $x \in [x_l, \alpha]$ so that setting $l = \alpha$ and $r = 2$ in line 9 is justified by the truth of $p_2(y, x_l, x_u, \alpha)$.

$y \preccurlyeq f(\alpha)$: in this case, setting $\mathrm{hi} = \alpha$ and $r = 1$ as done in line 12, is guaranteed to satisfy the post-condition: in fact, $n_g = 1$ implies $x_l \preccurlyeq \alpha \preccurlyeq \omega \preccurlyeq x_u$, therefore $\mathrm{hi} \in [x_l, x_u]$ and $f(\mathrm{hi}) \succcurlyeq y$.

Suppose now that the condition on line 2 is false: control flows to the else statement on line 14 and we know that either $n_\mathrm{g} > 1$ or $w_\mathrm{M} \leq t$. In this case function `linsearch_geq`$(f, y, [x_l, x_u], w_\mathrm{M}, t)$ searches backwards, float-by-float starting from x_u for at most the minimum between t and w_M steps, not going beyond x_l, the first value for hi such that $f(\mathrm{hi}) \succcurlyeq y$; b and hi as set in line 15 satisfy the

condition stated in lines 16–17. At this point, if $b = 1$, then a suitable value for hi was found and setting $r = 1$ satisfies the post-condition.

At line 20, since we are on the else statement, we know that $b = 0$. Therefore, by the post-condition stated in lines 16–17, we have that $\forall x \in [\hat{x}, x_u] : f(x) \prec y$, with $v = \min\{t, w_{\mathrm{M}}\}$ and $\hat{x} = \max\{x_l, {x_u}^{-^v}\}$. There are two cases:

- $t \geq w_{\mathrm{M}}$: in this case, since $w_{\mathrm{M}} \geq w^f_{\mathrm{M}}$, we have $t \geq w^f_{\mathrm{M}}$, therefore we are sure that the interval $[\hat{x}, x_u]$ contains the value for x just before the the last glitch began. We call such point x_m. By definition of glitch, we have that $\forall x \in [x_l, x_m] : f(x) \preccurlyeq f(x_m)$. Since $x_m \in [\hat{x}, x_u]$ we have that $f(x_m) \prec y$, and therefore, $\forall x \in [x_l, x_m] : f(x) \preccurlyeq f(x_m) \prec y$. Together with $\forall x \in [\hat{x}, x_u] : f(x) \prec y$ we have that $\forall x \in [x_l, x_u] : f(x) \prec y$ and setting $r = 0$ in line 21 satisfies the post-condition.
- $t < w_{\mathrm{M}}$: in this case, we cannot be sure that there are not solutions in interval $[x_l, x_u]$ but we could not find (with a complexity bounded by t) a suitable value for hi to start bisection. Hence, setting $l = x_l$ and $r = 2$ in line 23 satisfies, trivially, the post-condition.

□

Lemma 3. *Function* `bisect_lb` *satisfies its contract.*

Proof. Assume that the precondition for `bisect_lb`$(f, y, [x_l, x_u], n_{\mathrm{g}}, d_{\mathrm{M}}, w_{\mathrm{M}}, \alpha, \omega, n_g, s, t)$ is satisfied initially, and consider the following **while**-loop invariant: $Inv = (x_l \preccurlyeq \mathrm{lo} \prec \mathrm{hi} \preccurlyeq x_u) \wedge (f(\mathrm{lo}) \prec y \preccurlyeq f(\mathrm{hi})) \wedge (\forall x \in [x_l, \mathrm{lo}] : f(x) \prec y)$. We will show that when the **while**-loop of `bisect_lb` finishes, the post-condition in the **Ensure** statement holds. To this aim, we will prove that the **while**-loop of `bisect_lb` satisfies the following properties.

Initialization: the invariant Inv holds prior to the first loop iteration because it is entailed by the precondition, that is, the **Require** statement.

Maintenance: assume that Inv holds at the beginning of an arbitrary loop iteration: we will prove that it holds at the end of that iteration, as well. If the guard on line 1, $\mathrm{hi} \succ^1 \mathrm{lo}$, is false, than Inv trivially holds at the end of the loop. Therefore, assuming that $\mathrm{hi} \succ^1 \mathrm{lo}$ holds, a stronger property can be proved: Inv holds and either the loop is exited with a break statement, or the new value hi' or lo', that is updated during the iteration, is contained into the open interval $(\mathrm{lo}, \mathrm{hi})$.

After the invocation of function `split_point`, at line 3, $\mathrm{mid} = \mathrm{hi}^{-^m} = \mathrm{lo}^{+^{m'}}$, with $m, m' > 0$, holds. Hence, $\mathrm{lo} \prec \mathrm{mid} \prec \mathrm{hi}$ holds and this property is exploited in the following two main cases, tested at line 4.

- $y \preccurlyeq f(\mathrm{mid})$: in this case $\mathrm{hi}' = \mathrm{mid}$. By assuming that Inv holds before the iteration, we need to prove that $(x_l \preccurlyeq \mathrm{lo} \prec \mathrm{hi}' \prec \mathrm{hi} \preccurlyeq x_u) \wedge (f(\mathrm{lo}) \prec y \preccurlyeq f(\mathrm{hi}'))$. Note that this is a direct consequence of $\mathrm{lo} \prec \mathrm{mid} \prec \mathrm{hi}$ and $y \preccurlyeq f(\mathrm{mid})$. Since the value of lo was not updated, the remaining part of invariant Inv holds as it did before this iteration. Therefore, in the rest of the proof we will show that $\forall x \in [x_l, \mathrm{lo}] : f(x) \prec y$ only when the value of lo is being updated.

$y \succ f(\text{mid})$: here, further cases need to be distinguished.

$n_{\text{g}} = 0$ or $\text{mid} \preccurlyeq \alpha$ or $\text{mid} \succcurlyeq \omega$ or $y \succ^{d_{\text{M}}} f(\text{mid})$: in this case $\text{lo}' = \text{mid}$. From $\text{lo} \prec \text{mid} \prec \text{hi}$ and $y \succ f(\text{mid})$, we can derive that $(x_l \preccurlyeq \text{lo} \prec \text{lo}' \prec \text{hi} \preccurlyeq x_u) \wedge (f(\text{lo}') \prec y \preccurlyeq f(\text{hi}))$. We are left to prove that $\forall x \in [x_l, \text{lo}'] : f(x) \prec y$. If the function f is isotonic or, at least, we are sure that point mid is not inside a glitch, that is, $\text{mid} \preccurlyeq \alpha \preccurlyeq \alpha^f$ or $\text{mid} \succcurlyeq \omega \succcurlyeq \omega^f$, $y \succ f(\text{mid})$ implies that $\forall x \in [x_l, \text{lo}' = \text{mid}] : f(x) \prec y$ and the invariant Inv is proved. For the remaining case $y \succ^{d_{\text{M}}} f(\text{mid})$, note that $f(\text{mid})$ is so low that, even under the worst-case assumption mid is the minimum of a maximal-depth glitch, the value of f just before the glitch would still be smaller than y. Therefore, we can conclude that $\forall x \in [x_l, \text{lo}' = \text{mid}] : f(x) \prec y$.

$n_{\text{g}} = 1$ but either $w_{\text{M}} > t$ or $f(\alpha^+) \prec f(\alpha)$, and the previous condition is false: since line 8 is the **else-if** guard of the **if** statement at line 6, we are sure that $\alpha \prec \text{mid} \prec \omega$. This implies that mid *may be inside a glitch*. For this reason, at line 9 the condition $f(\omega) \succcurlyeq y$ is tested and if it holds, at line 10 condition $f(\alpha) \succcurlyeq y$ is tested. If it holds, we set $\text{hi} = \alpha$. Since $\text{mid} \prec \text{hi} \preccurlyeq x_u$ and $\alpha \prec \text{mid}$, we can conclude that $\alpha \prec \text{hi}$. Moreover, since $f(\text{lo}) \prec y$, $f(\alpha) \succcurlyeq y$ and $\alpha \preccurlyeq \alpha^f$, we have that $\text{lo} \prec \alpha$. Hence, we can conclude $(x_l \preccurlyeq \text{lo} \prec \text{hi}' = \alpha \prec \text{hi} \preccurlyeq x_u) \wedge (f(\text{lo}) \prec y \preccurlyeq f(\text{hi}'))$ holds.

At line 12 we know that $f(\alpha) \prec y \preccurlyeq f(\omega)$. Then we test if $f(\alpha^+) \prec f(\alpha)$: if this is true, we can be sure that the unique glitch of function f starts exactly at point α, that is $\alpha = \alpha^f$. Therefore, we have $f(\alpha^f) \prec y \preccurlyeq f(\omega)$. In this case $\text{lo}' = \text{mid}$. From $\text{lo} \prec \text{mid} \prec \text{hi}$ and $y \succ f(\text{mid})$, we can derive that $(x_l \preccurlyeq \text{lo} \prec \text{lo}' \prec \text{hi} \preccurlyeq x_u) \wedge (f(\text{lo}') \prec y \preccurlyeq f(\text{hi}))$. Moreover, $f(\alpha^f) \prec y \preccurlyeq f(\omega)$ and $y \succ f(\text{mid})$ implies that $\forall x \in [x_l, \text{lo}' = \text{mid}] : f(x) \prec y$: the last part of the invariant Inv is proved.

At line 14 $f(\alpha) \prec y \preccurlyeq f(\omega)$ still holds. In this case, if $\text{lo} \prec \alpha$, we set $\text{lo}' = \alpha$. From $\text{lo} \prec \alpha$, $\alpha \preccurlyeq \alpha^f$, $f(\alpha) \prec y$ and $y \preccurlyeq f(\text{hi})$ we can derive that $(x_l \preccurlyeq \text{lo} \prec \text{lo}' = \alpha \prec \text{hi} \preccurlyeq x_u) \wedge (f(\text{lo}') \prec y \preccurlyeq f(\text{hi}))$ and $\forall x \in [x_l, \text{lo}' = \alpha] : f(x) \prec y$, since $f(\alpha) \prec y$ and $\alpha \preccurlyeq \alpha^f$. This proves the invariant Inv.

At line 16 we know that $\text{lo} = \alpha$ and we exit from the loop. Note that, in this case, the invariant Inv it is trivially proved.

At line 19 we are sure that $f(\omega) \prec y$ and we set $\text{lo}' = \omega$. Since $f(\omega) \prec y$, $\omega \succcurlyeq \omega^f$ and $y \preccurlyeq f(\text{hi})$, we can conclude that $\omega \prec \text{hi} \preccurlyeq x_u$. Moreover, observe that $\omega \preccurlyeq \text{lo}$ cannot hold: since the **else-if**-guard of line 6 is false here, $\text{mid} \preccurlyeq \omega$. This would imply $\text{mid} \preccurlyeq \text{lo}$, which contradicts the post-condition of function `split_point`. Therefore, $x_l \preccurlyeq \text{lo} \prec \omega \prec \text{hi} \preccurlyeq x_u$. Moreover, by $f(\omega) \prec y$ and $\omega \succcurlyeq \omega^f$ we have $\forall x \in [x_l, \text{lo}' = \omega] : f(x) \prec y$. Setting $\text{lo} := \omega$ satisfies the invariant.

$w_{\text{M}} \leq t$ and the previous conditions are false: at line 22 function $\texttt{findfmax}(f, w_{\text{M}}, \text{lo}, \text{mid})$ is called, returning a value b that satisfies

the post-condition at line 23. The **if** statement of line 24 distinguishes between two cases:

$f(b) \succcurlyeq y$: $\mathrm{hi}' = b$ is set. Note that by the post-condition at line 23 we have that $b \in [\max\{\mathrm{lo}, \mathrm{mid}^{-^{w_\mathrm{M}}}\}, \mathrm{mid}]$. Since $\mathrm{lo} \prec \mathrm{mid} \prec \mathrm{hi}$ and $f(b) \succcurlyeq y$ while $f(\mathrm{lo}) \prec y$, we have that $\mathrm{lo} \prec b \prec \mathrm{hi}$. Therefore, $x_l \preccurlyeq \mathrm{lo} \prec \mathrm{hi}' = b \prec \mathrm{hi} \preccurlyeq x_u$.

$f(b) \prec y$: $\mathrm{lo}' = \mathrm{mid}$ is set. Since $\mathrm{lo} \prec \mathrm{mid} \prec \mathrm{hi}$ we have that $x_l \preccurlyeq \mathrm{lo} \prec \mathrm{lo}' = \mathrm{mid} \prec \mathrm{hi} \preccurlyeq x_u$. Moreover, $f(\mathrm{lo}') \prec y \preccurlyeq f(\mathrm{hi})$, since line 27 is in the **else** body of the **if** construct of line 4. Finally, we have to prove that $\forall x \in [x_l, \mathrm{lo}'] : f(x) \prec y$. From post-condition at line 23 we know that $\forall x \in [\max\{\mathrm{lo}, \mathrm{mid}^{-^{w_\mathrm{M}}}\}, \mathrm{mid}] : f(x) \preccurlyeq f(b)$. We have now two cases:

$\max\{\mathrm{lo}, \mathrm{mid}^{-^{w_\mathrm{M}}}\} = \mathrm{lo}$: In this case, post-condition at line 23 implies that $\forall x \in [\mathrm{lo}, \mathrm{mid}] : f(x) \preccurlyeq f(b) \prec y$.

$\max\{\mathrm{lo}, \mathrm{mid}^{-^{w_\mathrm{M}}}\} = \mathrm{mid}^{-^{w_\mathrm{M}}}$: Since $w_\mathrm{M} \geq w_\mathrm{M}^f$, even in the worst case scenario, that is, mid is inside a glitch at the maximal distance from the its beginning, interval $[\mathrm{mid}^{-^{w_\mathrm{M}}}, \mathrm{mid}]$ contains the last point x_m before the glitch started. Therefore, we have that $\forall x \in [\mathrm{lo}, x_m] : f(x) \preccurlyeq f(x_m)$. Together with $\forall x \in [\mathrm{mid}^{-^{w_\mathrm{M}}}, \mathrm{mid}] : f(x) \preccurlyeq f(b) \prec y$, we have $\forall x \in [\mathrm{lo}, \mathrm{mid}] : f(x) \preccurlyeq f(b) \prec y$, which proves that $\forall x \in [x_l, \mathrm{lo}'] : f(x) \prec y$.

In both cases, invariant Inv holds.

Otherwise: Function `logsearch_lb`$(f, d_\mathrm{M}, \mathrm{lo}, \mathrm{mid}, y, s)$, which returns a value z satisfying the post-condition stated at line 31, is called at line 30. The if statement of line 32 distinguishes between two cases:

$\mathrm{lo} \prec z$: We set $\mathrm{lo}' = z$. From post-condition stated at line 31, we have that $z \in [\mathrm{lo}, \mathrm{mid}]$. Together with $\mathrm{lo} \prec z$, this gives $x_l \preccurlyeq \mathrm{lo} \prec \mathrm{lo}' = z \prec \mathrm{hi} \preccurlyeq x_u$. Moreover, post-condition stated at line 31 implies that $f(z) \prec^{d_\mathrm{M}} y$, therefore also $f(\mathrm{lo}') \prec y \preccurlyeq f(\mathrm{hi}))$ holds. Finally, we have to prove that $\forall x \in [x_l, \mathrm{lo}' = z] : f(x) \prec y$. Since $f(z) \prec^{d_\mathrm{M}} y$ this means that $f(z)$ is so low that, even under the worst-case assumption z is the minimum of a maximal-depth glitch, the value of f just before the glitch would still be below y. This implies that for each x in the interval $[x_l, \mathrm{lo}' = z] : f(x) \prec y$.

$\mathrm{lo} = z$: in this case we exit the loop without any change and, therefore, invariant Inv trivially holds.

Termination: We have just proved that invariant Inv holds and either we exit the loop with a **break** statement, or the new value hi' or lo' is contained into the interval $(\mathrm{lo}, \mathrm{hi})$. Therefore, $\#[\mathrm{lo}, \mathrm{hi}]$ decreases at each iteration. The guard of the **while** loop at line 1 tests the condition $\mathrm{hi} \succ^1 \mathrm{lo}$, that is equivalent to $\#[\mathrm{lo}, \mathrm{hi}] > 2$. This assures that the **while** loop always terminates.

Correctness: We will prove that, whenever the loop invariant and the loop exit-condition both hold, then the post-condition stated in the **Ensure** statement holds. Since the "Correctness" post-condition coincides with the invariant Inv, we only need to prove that the "Precision" post-condition holds. Note

that, under the hypothesis that $n_{\mathrm{g}} = 0$, $w_{\mathrm{M}} < t$ or $(n_{\mathrm{g}} = 1 \wedge \alpha = \alpha^f)$, the control flow of the program can never reach the **break** statements at lines 17 or 35. If we did not exit with one of the above mentioned **break** statements, $\#[\mathrm{lo}, \mathrm{hi}] = 2$ holds, that is $\mathrm{lo}^+ = \mathrm{hi}$. Therefore, we can conclude that in these cases when we exit the loop $\mathrm{lo}^+ = \mathrm{hi}$ holds. Since for the Invariant Inv we have that $y \preccurlyeq f(\mathrm{hi})$, we can conclude that $f(\mathrm{lo}^+) \succcurlyeq y$. This concludes the proof.

□

Theorem 1. *Function* `lower_bound` *satisfies its contract.*

Proof. Assume that the precondition for $\texttt{lower_bound}(y, [x_l, x_u], n_{\mathrm{g}},\ d_{\mathrm{M}}, w_{\mathrm{M}}, \alpha, \omega, f^{\mathrm{i}}, t)$ is satisfied initially. First, the "Correctness" post-condition will be proved. As described in Section 5.2, function $\texttt{init}(y, [x_l, x_u], f^{\mathrm{i}})$ has been designed to return a point inside the interval $[x_l, x_u]$, which, therefore, meets the condition at line 2. Moreover, recall that function $\texttt{gallop_lb}(f, y, [x_l, x_u], d_{\mathrm{M}}, i)$ starts with $\mathrm{lo} = \mathrm{hi} = i$, where $x_l \preccurlyeq i \preccurlyeq x_u$ and it is realized in such a way that it returns lo and hi satisfying the condition at line 4. The goal now is to verify if such lo and hi can be used for bisection.

We first focus on the value of lo: the value $f(\mathrm{lo})$ is compared to y on line 5 and on line 11.

$f(\mathrm{lo}) \succ y$: In this case, using the post-condition at line 4, we can be sure that $\mathrm{lo} = x_l$. Since $f(x_l) \succ y$, we further need to distinguish the case in which we are guaranteed that $\forall x \in [x_l, x_u] : f(x) \succ y$ from the one we are not. The first case surely occurs in the following conditions, which are part of the guard of the **if** statement on line 6:

- $n_{\mathrm{g}} = 0$: in this case, since $n_{\mathrm{g}} \geq n_{\mathrm{g}}^f$, we can conclude that $n_{\mathrm{g}}^f = 0$ and, therefore, that function f is isotonic on the interval $[x_l, x_u]$. Since $f(\mathrm{lo}) \succ y$ we can be sure that $\forall x \in [x_l, x_u] : f(x) \succ y$.
- $f(\alpha) \succ^{d_{\mathrm{M}}} y$: Since $d_{\mathrm{M}} \geq d_{\mathrm{M}}^f$ from the precondition, we can conclude that $f(\alpha) \succ^{d_{\mathrm{M}}^f} y$. Moreover, $\alpha \preccurlyeq \alpha^f$ and the fact that f is quasi-isotonic allow us to conclude that $f(\alpha^f) \succcurlyeq f(\alpha) \succ^{d_{\mathrm{M}}^f} y$. By definition of α^f we know that function f is isotonic on interval $[x_l, \alpha^f]$, and since $f(x_l) \succ y$, we can conclude that $\forall x \in [x_l, \alpha^f]$, $f(x) \succ y$. Moreover, $f(\alpha^f) \succ^{d_{\mathrm{M}}^f} y$ assures us that the value $f(\alpha^f)$ is so high that, even under the worst case assumption that the glitch starting at α^f has maximal depth, the value of f could not decrease enough to reach y. Therefore, we can be sure that $\forall x \in [x_l, x_u] : f(x) \succ y$.

Hence, in all circumstances in which the condition on line 6 is true, the post-condition for the case $r = 1$ and $l = x_u$ is proved.
At line 8, we are on the **else** statement, therefore we cannot conclude that $\forall x \in [x_l, x_u] : f(x) \succ y$, but we cannot apply bisection, since we are under the assumption $f(\mathrm{lo}) \succ y$. From $\alpha \preccurlyeq \alpha^f$ and $f(x_l) \succ y$, we have that $\forall x \in [x_l, \alpha]$, $f(x) \succ y$. Therefore, setting $l = \alpha \prec x_u$ satisfies the post-condition for the case $r = 1$.

$f(\text{lo}) = y$: In this case, using the condition at line 4, we can be sure that $x_l = \text{lo}$. Therefore, setting $l = \text{lo}$ satisfies the post-condition for the case $r = 4$.

Since each branch of the **if-else** instructions considered until now contains a return statement, at the end of line 13 we are sure that $f(\text{lo}) \prec y$.

At line 14, we focus on the value of hi. If $f(\text{hi}) \prec y$, using the condition at line 4 we are sure that $\text{hi} = x_u$ and, hence, $f(x_u) \prec y$. Under this condition, the value of $\text{hi} = x_u$ cannot can be used for bisection: function `findhi_lb`$(f, y, [x_l, x_u], n_\text{g}, d_\text{M}, w_\text{M}, \alpha, \omega, t)$ specified in Algorithm 2 is called. The precondition of function `findhi_lb` is satisfied by the one of function `lower_bound` and by the hypothesis $f(\text{hi}) \prec y$. Afterwards, by Lemma 2, the post-condition of function `findhi_lb` holds. At line 16, we test if function `findhi_lb` returned $r = 0$ or $r = 2$: in these cases the post-condition of function `findhi_lb` assures that the one of function `lower_bound` for $r = 0$ and $r = 2$ is satisfied. At the end of line 17, we know that function `findhi_lb` returned $r = 1$ and, therefore, a new value such that $x_l \preccurlyeq \text{hi} \preccurlyeq x_u$ and $f(\text{hi}) \succcurlyeq y$ was assigned to hi.

Summarizing, before the invocation of function `bisect_lb` at line 19 we are sure that $f(\text{lo}) \prec y \preccurlyeq f(\text{hi})$ and, therefore, $\text{lo} \neq \text{hi}$.

In order to prove that the preconditions of function `bisect_lb` are met, we need to prove that the values of lo and hi at line 19 satisfy the following conditions: $x_l \preccurlyeq \text{lo} \prec \text{hi} \preccurlyeq x_u$ and $\forall x \in [x_l, \text{lo}] : f(x) \prec y$. Note, indeed, that we have already proved that at line 19 $f(\text{lo}) \prec y \preccurlyeq f(\text{hi})$ holds and that the remaining requirements of unction `bisect_lb` are implied by the precondition of function `lower_bound`. Moreover, we will first prove that the value of lo at line 19 satisfies $x_l \preccurlyeq \text{lo}$ and $\forall x \in [x_l, \text{lo}] : f(x) \prec y$. Observe that the value of lo is the one returned by function `gallop_lb`, since the code at lines 5-19 tests the value of lo without changing it. Therefore, by the post-condition of line 4, we have that $x_l \preccurlyeq \text{lo}$. As for the remaining clause, we have two cases:

- $x_l \prec \text{lo}$: by the post-condition of line 4, we have that $y \succ^{d_\text{M}} f(\text{lo})$. This means that $f(\text{lo})$ is so low that, even under the worst-case assumption lo is the minimum of a maximal-depth glitch, the value of f just before the glitch would still be below y. This implies that $\forall x \in [x_l, \text{lo}] : f(x) \prec y$.
- $x_l = \text{lo}$: since at line 19 we are sure that $f(x_l) \prec y$, in this case $\forall x \in [x_l, \text{lo}] : f(x) \prec y$ holds, trivially.

In order to prove that the value of hi at line 19 satisfies $\text{lo} \prec \text{hi} \preccurlyeq x_u$, we need to distinguish between the following cases:

- $y \preccurlyeq f(\text{hi})$: the value of hi is the one computed by `gallop_lb`, since the guard of the **if**-instruction at line 14 is false. Therefore, by the condition of line 4, we have that $\text{lo} \preccurlyeq \text{hi} \preccurlyeq x_u$. Since at line 19 we are sure that $\text{lo} \neq \text{hi}$, we can conclude that $\text{lo} \prec \text{hi} \preccurlyeq x_u$.
- $y \succ f(\text{hi})$: in this case, the current value of hi is the one chosen by function `findhi_lb`, which must have returned $r = 1$, since we are at line 19. Therefore, by the post-condition of function `findhi_lb`, we have that the new value of hi is such that $\text{hi} \in [x_l, x_u]$ and $f(\text{hi}) \succcurlyeq y$. Since we have proved that $\forall x \in [x_l, \text{lo}] : f(x) \prec y$, we can conclude that $\text{lo} \prec \text{hi} \preccurlyeq x_u$.

Hence, since the preconditions of function `bisect_lb` are met, by Lemma 3 we are sure that, after it returns, $x_l \preccurlyeq \text{lo} \prec \text{hi} \preccurlyeq x_u$, $f(\text{lo}) \prec y \preccurlyeq f(\text{hi})$ and $\forall x \in [x_l, \text{lo}] : f(x) \prec y$ hold, for the new values of lo and hi. At line 20 we enter in a **while**-loop. This loop realizes a float-by-float search (for at maximum t iterations) to approach the exact solution of $y = f(\mathbf{x})$. We want to prove at line 23 the following property, that is also the invariant of the loop: $\forall x \in [x_l, \text{lo}] : f(x) \prec y$. To this aim, we will prove the following properties of the loop:

Initialization: the invariant $\forall x \in [x_l, \text{lo}] : f(x) \prec y$ holds prior to the first loop iteration because it is entailed by the post-condition of function `bisect_lb`.

Maintenance: assume $\forall x \in [x_l, \text{lo}] : f(x) \prec y$ holds at the beginning of an arbitrary loop iteration. Then such assumption, together with the guard of the loop $f(\text{lo}^+) \prec y$ and the assignment $\text{lo}' := \text{lo}^+$ in the body of the loop, allows us to conclude that $\forall x \in [x_l, \text{lo}'] : f(x) \prec y$ holds also at the end of that iteration.

Termination: the **while** loop terminates because we are assured by the post-condition of `bisect_lb` that there exists a value hi such that $\text{lo} \prec \text{hi}$ and $y \preccurlyeq f(\text{hi})$. Moreover, the loop can end before reaching such value, because the parameter $t \in \mathbb{N}$ is decremented inside the loop, until it reaches 0.

Correctness: as a consequence, at the end of the **while**-loop the property $\forall x \in [x_l, \text{lo}] : f(x) \prec y$ holds and either $f(\text{lo}^+) \succcurlyeq y$ or $t = 0$.

Then, at line 24 we test if $f(\text{lo}^+) \succ y$. Since $\forall x \in [x_l, \text{lo}] : f(x) \prec y$, setting l to the value of lo satisfies the post-condition for the case $r = 3$. With the **else-if** instruction at line 26 we test if $f(\text{lo}^+) = y$. In this case, since $\forall x \in [x_l, \text{lo}] : f(x) \prec y$, setting l to the value of lo^+ satisfies the post-condition for the case $r = 4$. Furthermore, if we are on the **else** instruction of line 28, it means that the **while**-loop of line 20 terminated because t became equal to 0. In this case, since $\forall x \in [x_l, \text{lo}] : f(x) \prec y$, setting l to the value of lo satisfies the post-condition for the case $r = 2$.

We will now prove the "Precision" post-condition under the hypothesis that $f(x_l) \preccurlyeq y \preccurlyeq f(x_u)$ holds and either $n_\text{g} = 0$, $w_\text{M} < t$ or $n_\text{g} = 1 \wedge \alpha = \alpha^f$ holds. Reasoning as before, at line 3 we are sure that the function `gallop_lb`$(f, y, [x_l, x_u], d_\text{M}, i)$ returns values for lo and hi such that post-condition at line 4 is satisfied. The latter, together with the assumption $f(x_l) \preccurlyeq y \preccurlyeq f(x_u)$, allows us to conclude that, in the cases we are interested in, $f(\text{lo}) \preccurlyeq y \preccurlyeq f(\text{hi})$. The condition tested by the if statement at line 5 is surely false. Therefore, in this case, condition at line 11 needs to be tested. We have the following cases:

- $f(\text{lo}) = y$: function `lower_bound` returns with $r = 4$. Therefore, the "Precision" post-condition is satisfied.
- $f(\text{lo}) \prec y$: since $y \preccurlyeq f(\text{hi})$, the test of line 14 is false and before line 19 we are sure that $x_l \preccurlyeq \text{lo} \prec \text{hi} \preccurlyeq x_u$, since $f(x_l) \prec y$ and $y \preccurlyeq f(x_u)$. Therefore, since the preconditions of function `bisect_lb` are met, by Lemma 3 we have that the post-condition holds. In more details, since we have assumed that $n_\text{g} = 0 \vee w_\text{M} < t \vee (n_\text{g} = 1 \wedge \alpha = \alpha^f)$, from the "Precision" post-condition of

function `bisect_lb` we can derive that $f(\mathrm{lo}^+) \succcurlyeq y$. In this case, the **while**-loop of line 20 is never executed. The value of $f(\mathrm{lo}^+) \succcurlyeq y$ is then tested at lines 24 and 26. If $f(\mathrm{lo}^+) \succ y$ then $r = 3$ is returned, if $f(\mathrm{lo}^+) = y$, then $r = 4$ is returned. In both cases, the "Precision" post-condition is satisfied.

This completes the proof. □

Theorem 2. *In case $f\colon \mathbb{F} \to \mathbb{F}$ is isotonic, that is $n_\mathrm{g} = 0$, then, for any $[x_l, x_u] \in \mathcal{I}_\mathbb{F}$, d_M, w_M, α, ω, $f^\mathrm{i}\colon \mathbb{F} \to \mathbb{F}$, $s, t \in \mathbb{N}$, the call to function* `lower_bound` $(f, y, [x_l, x_u], 0, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, f^\mathrm{i}, s, t)$ *in Algorithm 1 can be executed calling at most $2\log_2(\#[x_l, x_u]) + 4$ times function f.*

Proof. At the beginning of Algorithm 1, function $\texttt{init}(y, [x_l, x_u], f^\mathrm{i})$ is called and it returns a point inside interval $[x_l, x_u]$, meeting the condition at line 2. It never calls function f. Then, function $\texttt{gallop_lb}(f, y, [x_l, x_u], d_\mathrm{M}, i)$ starts with $\mathrm{lo} = \mathrm{hi} = i$. First, assume that $d_\mathrm{M} = 0$. Since either $y \succ^0 f(i)$ (that is, $y \succ f(i)$) or $f(i) \succcurlyeq y$ holds, the worst cases are when:

- $i = x_l$ and $\forall x \in [x_l, x_u) : y \succ f(x)$: in this case we find $\mathrm{hi} = x_u$ after $\log_2(\#[x_l, x_u]) + 1$ calls to function f;
- $i = x_u$ and $\forall x \in (x_l, x_u] : f(x) \succcurlyeq y$: in this case we find $\mathrm{lo} = x_l$ after $\log_2(\#[x_l, x_u]) + 1$ calls to f.

For the case $d_\mathrm{M} > 0$, note that the worst case for the function `gallop_lb` is when it starts at some point inside interval $[x_l, x_u]$, but it terminates only when $\mathrm{lo} = x_l$ and $\mathrm{hi} = x_u$. In this case `gallop_lb` invokes $\log_2(\#[x_l, x_u]) + 1$ times function f, as well.

After calling `gallop_lb`, we have the following cases.

- $f(\mathrm{lo}) \succcurlyeq y$: condition $f(\mathrm{lo}) \succ y$ is tested at line 5 of Algorithm 1. Then, since $n_\mathrm{g} = 0$ the function returns at line 7. Condition $f(\mathrm{lo}) = y$ is tested at line 11 and the function returns at line 12.
- $f(\mathrm{hi}) \prec y$: function $\texttt{findhi_lb}(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, t)$ is called but, since $n_\mathrm{g} = 0$, $r = 0$ is returned to the calling function `lower_bound`, that ends at line 16 without invoking function f.
- $f(\mathrm{lo}) \prec y \preccurlyeq f(\mathrm{hi})$: at line 19 we call function $\texttt{bisect_lb}(f, y, n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, s, t, \mathrm{lo}, \mathrm{hi})$. At each iteration, it computes the middle point mid between lo and hi, and $f(\mathrm{mid})$. Since $n_\mathrm{g} = 0$, it updates either the value of hi or the value of lo with mid until $\mathrm{lo}^+ = \mathrm{hi}$. Since mid satisfies the post-condition of function `split_point` at line 3 of Algorithm 3, we can conclude that $\#[\mathrm{mid}, \mathrm{hi}] \leq \lceil \#[\mathrm{lo}, \mathrm{hi}]/2 \rceil$ and $\#[\mathrm{lo}, \mathrm{mid}] \leq \lceil \#[\mathrm{lo}, \mathrm{hi}]/2 \rceil$. Hence, in the worst case function `bisect_lb` calls f $\log_2(\#[x_l, x_u])$ times. When $n_\mathrm{g} = 0$ the "Precision" post-condition of `bisect_lb` holds, then we are sure that $f(\mathrm{lo}^+) \succcurlyeq y$. At line 20 of function `lower_bound`, we find a **while**-loop. Since $\mathrm{lo}^+ = \mathrm{hi}$ and $f(\mathrm{lo}^+) \succcurlyeq y$, the **while**-loop of line 20 is never executed. Then, either the function `lower_bound` returns at line 25 because $f(\mathrm{lo}^+) \succ y$, or it returns at line 27 because $f(\mathrm{lo}^+) = y$. In any case, function `lower_bound` called function f at most $2\log_2(\#[x_l, x_u]) + 4$ times.

This concludes the proof. □

Theorem 3. *In case the function* $f\colon \mathbb{F} \to \mathbb{F}$ *has small glitches, that is* $n_{\mathrm{g}} > 0$ *but* $w_{\mathrm{M}} < t$*, then, for any* $[x_l, x_u] \in \mathcal{I}_{\mathbb{F}}$*,* d_{M}*,* α*,* ω*,* $f^{\mathrm{i}}\colon \mathbb{F} \to \mathbb{F}$*,* $s \in \mathbb{N}$*, the call to function* `lower_bound`$(f, y, [x_l, x_u], n_{\mathrm{g}}, d_{\mathrm{M}}, w_{\mathrm{M}}, \alpha, \omega, f^i, s, t)$ *in* Algorithm 1 *can be executed calling at most* $(w_{\mathrm{M}}+1)\log_2(\#[x_l, x_u]) + 6 + w_{\mathrm{M}}$ *times the function* f*.*

Proof. Algorithm 1, describing function `lower_bound` starts by calling function `init`$(y, [x_l, x_u], f^{\mathrm{i}})$, which returns a point inside interval $[x_l, x_u]$ without calling function f. Then, at line 3 of Algorithm 1, function `gallop_lb`$(f, y, [x_l, x_u], d_{\mathrm{M}}, i)$ is invoked. As we discussed in the proof of Theorem 2, in the worst case, `gallop_lb` invokes $\log_2(\#[x_l, x_u]) + 1$ times function f.

Then we have the two following cases.

- $f(\mathrm{lo}) \succcurlyeq y$: condition $f(\mathrm{lo}) \succ y$ is tested at line 5. Then, since $n_{\mathrm{g}} > 0$, function `lower_bound` of Algorithm 1 returns either on line 7 or on line 9, depending on whether $f(\alpha) \succ^{d_{\mathrm{M}}} y$ holds or not. If $f(\mathrm{lo}) \preccurlyeq y$, condition $f(\mathrm{lo}) = y$ is tested at line 11 and if it holds, function `lower_bound` returns at line 12. Therefore, in this case function `lower_bound` terminates calling function f at most 2 times.
- $f(\mathrm{lo}) \preccurlyeq y$ but $f(\mathrm{hi}) \prec y$: function `findhi_lb`$(f, y, [x_l, x_u], n_{\mathrm{g}}, d_{\mathrm{M}}, w_{\mathrm{M}}, \alpha, \omega, t)$ is called after 2 calls to f (to test the values of $f(\mathrm{lo})$ and $f(\mathrm{hi})$). We will now prove that the call to `findhi_lb` will invoke function f at most $w_{\mathrm{M}} + 2$ times, in the worst case.

  Since $n_{\mathrm{g}} > 0$ and $w_{\mathrm{M}} < t$, we have the following cases regarding the function `findhi_lb`.
  - $x_u \succ \omega$ or $y \succ^{d_{\mathrm{M}}} f(x_u)$: the call to `findhi_lb` terminates at line 3, returning the value $r = 0$. This causes function `lower_bound` to terminate with 2 invocations to function f.
  - $n_{\mathrm{g}} = 1$, $f(\alpha^+) \prec f(\alpha)$ and $y \succ f(\alpha)$: the call to `findhi_lb` terminates at line 7 returning the value $r = 0$. Again, `lower_bound` terminates with 2 calls to function f.
  - $n_{\mathrm{g}} = 1$, $f(\alpha^+) \prec f(\alpha)$ and $y \preccurlyeq f(\alpha)$: the call to `findhi_lb` terminates at line 12 returning the value $r = 1$ and invoking the function f just 2 times.
  - $n_{\mathrm{g}} > 1$: after 2 invocations to function f, `findhi_lb` calls `linsearch_geq` $(f,\ y, [x_l, x_u], w_{\mathrm{M}}, t)$, which searches backwards for at most w_{M} steps (recall $w_{\mathrm{M}} < t$, in this case), float-by-float starting from $\mathrm{hi} = x_u$ the first value for hi such that $f(\mathrm{hi}) \succcurlyeq y$. Function `linsearch_geq` invokes the function f at most w_{M} times.

  Therefore we can conclude that `findhi_lb` will call function f at most $w_{\mathrm{M}}+2$ times. From a complexity point of view, the worst case is when `findhi_lb` returns a good hi to start bisection (in this case it returns $r = 1$, see lines 12 and 19 of Algorithm 2). Hence, at line 19 of Algorithm 1, we call function `bisect_lb`$(f, y, n_{\mathrm{g}}, d_{\mathrm{M}}, w_{\mathrm{M}}, \alpha, \omega, s, t, \mathrm{lo}, \mathrm{hi})$ that, at each iteration, computes the middle point mid between lo and hi and $f(\mathrm{mid})$. If $y \preccurlyeq f(\mathrm{mid})$, function

`bisect_lb` updates the value of hi with mid at line 5. As we have already discussed in the proof of Theorem 2, mid satisfies the post-condition of function `split_point` at line 3 of Algorithm 3. Therefore, $\#[\text{lo}, \text{mid}] \leq \lceil \#[\text{lo}, \text{hi}]/2 \rceil$. Alternatively, when $y \succ f(\text{mid})$, we have the following cases:

- mid $\preccurlyeq \alpha$ or mid $\succcurlyeq \omega$ or $y \succ^{d_{\text{M}}} f(\text{mid})$: function `bisect_lb` at line 7 updates the value of lo with mid. Since mid satisfies the post-condition of function `split_point` at line 3 of Algorithm 3, we have that $\#[\text{mid}, \text{hi}] \leq \lceil \#[\text{lo}, \text{hi}]/2 \rceil$.
- $\alpha \prec \text{mid} \prec \omega$, $n_{\text{g}} = 1$, $f(\alpha^+) \prec f(\alpha)$, $f(\omega) \succcurlyeq y$ and $f(\alpha) \succcurlyeq y$: at line 11 the value of hi is updated with α. Since $\alpha \prec \text{mid} \prec \omega$, we have that $\#[\text{lo}, \alpha] < \lceil \#[\text{lo}, \text{hi}]/2 \rceil$.
- $\alpha \prec \text{mid} \prec \omega$, $n_{\text{g}} = 1$, $f(\alpha^+) \prec f(\alpha)$, $f(\omega) \succcurlyeq y$ and $f(\alpha) \prec y$: at line 13 the value of lo is updated with mid. Remember that $\#[\text{mid}, \text{hi}] \leq \lceil \#[\text{lo}, \text{hi}]/2 \rceil$.
- $\alpha \prec \text{mid} \prec \omega$, $n_{\text{g}} = 1$, $f(\alpha^+) \prec f(\alpha)$ and $f(\omega) \prec y$: at line 19 the value of lo changes to ω. Since, in this case, $\alpha \prec \text{mid} \prec \omega$, we have that $\#[\omega, \text{hi}] < \lceil \#[\text{lo}, \text{hi}]/2 \rceil$
- $n_{\text{g}} > 1$: at line 22, function `findfmax`$(f, w_{\text{M}}, \text{lo}, \text{mid})$ is called. It finds a value b inside interval $[\max\{\text{lo}, \text{mid}^{-^{w_{\text{M}}}}\}, \text{mid}]$ for which the value $f(b)$ is the highest. Hence, function `findfmax` calls f w_M times. Then we have two cases based on the value b returned by `findfmax`:
  - $f(b) \succcurlyeq y$: at line 25, the value of hi is updated with b. Since $x_l \preccurlyeq b \preccurlyeq \text{mid}$, we have that $\#[\text{lo}, b] \leq \lceil \#[\text{lo}, \text{hi}]/2 \rceil$.
  - $f(b) \succ y$: at line 27 the value of lo is updated with mid and, as discussed above, $\#[\text{mid}, \text{hi}] \leq \lceil \#[\text{lo}, \text{hi}]/2 \rceil$ holds.

Hence, in the worst case function `bisect_lb` has to call function `findfmax` at each step. Therefore, in that case algorithm `bisect_lb` calls function f $w_M \log_2(\#[x_l, x_u])+1$ times. Since the "Precision" post-condition of `bisect_lb` holds ($w_{\text{M}} < t$), we are sure that $f(\text{lo}^+) \succcurlyeq y$. At line 20 of function `lower_bound`, there is a **while**-loop. Since $\text{lo}^+ = \text{hi}$ and $f(\text{lo}^+) \succcurlyeq y$, the said loop is never executed. Then, either function `lower_bound` returns at line 25 because $f(\text{lo}^+) \succ y$, or it returns at line 27 because $f(\text{lo}^+) = y$. In any case, function `lower_bound` called function f at most $(w_{\text{M}} + 1)\log_2(\#[x_l, x_u]) + 6 + w_{\text{M}}$ times.

This concludes the proof. □

B.2 Computation of Upper Bounds

Lemma 4. *Function* `findhi_ub` *specified in* Algorithm 5 *satisfies its contract.*

Proof. We assume that the precondition for `findhi_ub`$(f, y, [x_l, x_u], n_{\text{g}}, d_{\text{M}}, w_{\text{M}}, \alpha, \omega, t)$ is satisfied. In particular, $f(x_u) \prec y$ holds.

The guard of the **if** statement on line 1 is the same as the one of Algorithm 2. When it is satisfied we know for sure that either x_u is not in a glitch, or that x_u might be inside a glitch but not deep enough for the function to reach the value

y elsewhere in $[x_l, x_u]$. Together with $f(x_u) \prec y$, this allows us to state that the equation $y = f(\mathbf{x})$ has no solution, and setting $u = x_l$ satisfies p_6.

The purpose of the **if** block which starts at line 3 is to distinguish between a few cases when we know from $n_\mathrm{g} = 1$ and $n_\mathrm{g} \geq n_\mathrm{g}^f$ that f has at most one glitch. If the said block is entered, $x_u \preccurlyeq \omega$ and $y \preccurlyeq^{d_\mathrm{M}} f(x_u)$ hold, which comes from the negation of the previous **if** statement. Therefore, x_u may be inside a glitch. We compute the value of the function in α to discern between the cases listed below.

$y \prec f(\alpha)$: This, together with the precondition of this algorithm, implies that $f(x_u) \prec y \prec f(\alpha)$. Therefore, there must exist some $x \in [\alpha, x_u]$ such that $f(x) \preccurlyeq y$. If $w_\mathrm{M} > t$, then we refrain from searching such x because the glitch may be too wide, and setting $u = x_u$ trivially satisfies p_6. This is captured by the **if** guard of line 4. Otherwise, the control flow can reach line 12.

$y \succcurlyeq f(\alpha) \wedge f(\alpha^+) \succcurlyeq f(\alpha)$: The second term of this condition implies that $\alpha \prec \alpha^f$. Therefore, we cannot exclude that the function further increases between α and α^f, reaching y. If $w_\mathrm{M} > t$ there is nothing more we can do to refine the upper bound, and setting $u = x_u$ trivially satisfies p_6. This case is also captured by the **if** statement of line 4. If $w_\mathrm{M} \leq t$, line 12 can be executed.

$y = f(\alpha) \wedge f(\alpha^+) \prec f(\alpha)$: The second term of this condition implies that $\alpha = \alpha^f$. Since $y \succ f(x_u)$ and $n_\mathrm{g} = 1$, we know for sure that the whole interval $(\alpha, x_u]$ is a glitch, and the function is strictly lower than y in that interval. Therefore, we just found the rightmost point $x = \alpha$ where $f(x) = y$, and predicate p_6 is satisfied if $u = \alpha^+$ is set. When this condition is true, control flows into the **if** statement of line 6.

$y \succ f(\alpha) \wedge f(\alpha^+) \prec f(\alpha)$: The same reasoning we did for the previous case can be made: $\forall x \in [\alpha, x_u] : y \succ f(x)$. Together with the fact that $y \succ f(\alpha)$ and that f is isotonic to the right of α, this allows us to state that the equation $y = f(\mathbf{x})$ has no solutions in $[x_l, x_u]$, and we can set $u = x_l$, satisfying p_6. This last case is caught by the **else** statement of line 8.

The **if** guard of line 3 captures all the conditions listed above if $w_\mathrm{M} > t$. Otherwise, if $w_\mathrm{M} \leq t$, only the last two conditions are captured, because if the first two of them occur, then proceeding with the linear search described below can find a better value for u.

In line 12, function `linsearch_geq` is invoked. Starting from x_u, it searches backwards for a point where the functions reaches or exceeds y. If it does not find such value in at most t steps and before reaching x_l, $b = 0$ is set. Otherwise, $b = 1$ is set and hi is set to such value.

- If $b = 1$, the post-condition of `linsearch_geq` states that $\forall x \in [\mathrm{hi}^+, x_u] : y \succ f(x)$, and hi is the rightmost point where $y \preccurlyeq f(x)$. Therefore, setting $u = \mathrm{hi}^+$ satisfies p_6.
- If $b = 0$, then `linsearch_geq` could not find a point where the function is greater or equal than y, and the function is always lower than y in the interval $[\hat{x}, x_u]$. This could happen because of either of the two reasons described below, which can be distinguished by comparing t and w_M.

$t \geq w_{\mathrm{M}}$: `linsearch_geq` analyzed the entire glitch, and the value of f in the point where the glitch starts is still lower than y. Since the function is quasi-isotonic, it can only decrease further proceeding towards x_l. Therefore, we can state that the equation $y = f(\mathbf{x})$ has no solution, and set $u = x_l$, which satisfies p_6.

$t < w_{\mathrm{M}}$: This last case occurs when the **else** branch of line 20 is entered. The glitch was too wide to be analyzed completely, and we can only set $u = \hat{x}$, satisfying p_6.

□

Lemma 5. *Function* `check_glitch` *specified in* Algorithm 6 *satisfies its contract.*

Proof. The precondition of function `check_glitch` requires three values lo, m and hi such that $x_l \preccurlyeq \mathrm{lo} \preccurlyeq \mathrm{m} \preccurlyeq \mathrm{hi} \preccurlyeq x_u$ and both $f(\mathrm{m}) \succ y$ and $f(\mathrm{hi}) \succ y$. In this situation, we would like to claim that the graph of f is always higher than y, and that there is no solution to equation $y = f(\mathbf{x})$ in interval $[\mathrm{m}, \mathrm{hi}]$. However, there might be a glitch inside the said interval that lets the graph of function f reach the value y. The purpose of this algorithm is to search the said glitch float-by-float when possible, in order to either find the rightmost point where f evaluates to y, or to be able to claim that f never reaches y in the said interval.

The algorithm proceeds with the linear search only if there is at most one glitch, and it is sufficiently tight. This condition is checked by the guard of the **if** statement on line 1. If it does not hold, we refrain from searching the glitch because it is computationally too expensive, and b is set to 2 in line 13, which means that nothing can be said about the solution of the equation $y = f(\mathbf{x})$ in interval $[\mathrm{m}, \mathrm{hi}]$.

Otherwise, function `linsearch_leq` searches the glitch backwards starting from s_u for a maximum of w_M floats, until it either finds a value z such that $f(z) \preccurlyeq y$, or it reaches s_l. s_l and s_u are chosen so that the whole glitch is checked by `linsearch_leq`. Therefore, s_l is set to the maximum between $\alpha \preccurlyeq \alpha^f$ and lo. As for s_u, it is chosen depending on the following cases.

$f(\omega^-) \prec f(\omega)$: This implies that $\omega = \omega^f$, and the glitch finishes in ω. Setting s_u to the minimum between ω and hi lets `linsearch_leq` search the entire glitch (or the part of it inside the interval of interest), because it is surely not too wide ($w_M < t$).

$\alpha \prec^k \omega \wedge k \leq t$: We may not know exactly where the glitch ends, but the distance between α and ω is small enough to search the whole interval delimited by them. Again, setting $s_u := \min\{\omega, \mathrm{hi}\}$ lets `linsearch_leq` analyze the entire glitch.

$f(\alpha^+) \prec f(\alpha)$: This implies $\alpha = \alpha^f$. Therefore, the glitch starts exactly in α and, since its width is lower than t, setting $s_u := \min\{\alpha^{+w_{\mathrm{M}}}, \mathrm{hi}\}$ allows `linsearch_leq` to search the entire glitch.

On line 9 function `linsearch_leq` is finally called. It can only return $b = 0$ or $b = 1$, which proves part of the post-condition of this algorithm ($n_{\mathrm{g}} = 1 \wedge w_{\mathrm{M}} <$

$t \wedge \big(\alpha = \alpha^f \vee \omega = \omega^f \vee \big(\alpha \prec^k \omega \wedge k \leq t\big)\big) \implies b \in \{0,1\}$). It also satisfies the two other claims of the **Ensure** statement, depending on the value of b.

$b = 0$: By the post-condition of `linsearch_leq`, $\forall x \in [z, s_u] : f(x) \succ y$ holds, with $z = \hat{x} = \max\{s_l, s_u^{-w_\mathrm{M}}\}$. As we previously explained, s_l and s_u are chosen so that the only glitch is entirely contained in $[s_l, s_u]$, and it is within w_M floats from s_u: the glitch is contained in interval $[z, s_u]$, which was analyzed by `linsearch_leq`. Therefore, function f is actually isotonic in $[\mathrm{m}, z)$ and in $[s_u, \mathrm{hi}]$. This, together with $f(\mathrm{m}) \succ y$, $f(\alpha) \preccurlyeq f(\omega)$, and the post-condition of `linsearch_leq`, allows us to claim that $\forall x \in [\mathrm{m}, \mathrm{hi}] : f(x) \succ y$.

$b = 1$: The post-condition of `linsearch_leq` states that $z \in [\hat{x}, s_u]$, with $\hat{x} = \max\{s_l, s_u^{-w_\mathrm{M}}\} \succcurlyeq \mathrm{lo}$ and $s_u \preccurlyeq \mathrm{hi}$: this implies that $\mathrm{lo} \preccurlyeq z \preccurlyeq \mathrm{hi}$.
Moreover, s_u was chosen so that f is isotonic in the interval $[s_u, \mathrm{hi}]$. This, together with $\forall x \in (z, s_u] : f(x) \succ y$ (which is part of the post-condition of `linsearch_leq`), proves that $\forall x \in (z, \mathrm{hi}] : f(x) \succ y$.
Finally, the post-condition of `linsearch_leq` states that $f(z) \preccurlyeq y$, which concludes the proof of this part of the post-condition of `check_glitch`.

□

Lemma 6. *Function* `bisect_ub` *specified in* Algorithm 7 *satisfies its contract.*

Proof. We assume that the precondition of `bisect_ub` $(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, s, t, \mathrm{lo}, \mathrm{hi})$ is satisfied before the first iteration of the **while** loop that starts at line 1. We will now prove that the loop invariant $Inv = (x_l \preccurlyeq \mathrm{lo} \prec \mathrm{hi} \preccurlyeq x_u) \wedge (f(\mathrm{lo}) \preccurlyeq y \prec f(\mathrm{hi})) \wedge (\forall x \in [\mathrm{hi}, x_u] : f(x) \succ y)$ holds during and after the execution of the loop, satisfying the post-condition expressed in the **Ensure** statement.

Initialization: The invariant Inv is implied by the precondition of this algorithm. Therefore, it holds before the execution of the loop.

Maintenance: At the beginning of the loop body, we assume that both Inv and $\mathrm{hi} \succ^1 \mathrm{lo}$, the guard of the loop, hold. We will now prove that either the loop terminates with a **break** statement, or it continues and one between lo and hi takes a new value lo' or hi'. The said new value will be part of the interval $(\mathrm{lo}, \mathrm{hi})$. Note that this last statement implies $x_l \preccurlyeq \mathrm{lo} \preccurlyeq \mathrm{lo}' \preccurlyeq \mathrm{hi}' \preccurlyeq \mathrm{hi} \preccurlyeq x_u$, and therefore any new value for lo and hi that satisfies this condition, also satisfies the first condition of the invariant. We will prove that the invariant holds at the end of the loop body.
After function `split_point` is invoked at line 2, $\mathrm{lo} \prec \mathrm{mid} \prec \mathrm{hi}$. Note that this implies that, every time lo' or hi' are set to mid, the new value is part of the interval $(\mathrm{lo}, \mathrm{hi})$. Then, the value of $f(\mathrm{mid})$ is compared to y in order to decide whether mid can be a new value for lo or hi. The following cases may occur.

$f(\mathrm{mid}) \preccurlyeq y$: lo' is set to mid. Since $\mathrm{lo} \prec \mathrm{mid} \prec \mathrm{hi}$ and because Inv holds, also $x_l \preccurlyeq \mathrm{lo}' = \mathrm{mid} \prec \mathrm{hi} \preccurlyeq x_u$ holds. Again, Inv and $f(\mathrm{mid}) \preccurlyeq y$ imply that $f(\mathrm{lo}') \preccurlyeq y \prec f(\mathrm{hi})$. The third part of the invariant trivially holds, because hi remained the same.

$f(\text{mid}) \succ y$: If $n_g = 0 \vee \text{hi} \preccurlyeq \alpha \vee \text{mid} \succcurlyeq \omega \vee y \prec^{d_M} f(\text{mid})$ we are sure that $\forall x \in [\text{mid}, \text{hi}] : f(x) \succ y$, because there is no glitch in $[\text{mid}, \text{hi}]$ where the function could become lower than or equal to y. This can be stated because either we are sure there are no glitches at all in that interval (first three conditions of the **if** guard of line 5), or $f(\text{mid})$ is so higher than y that no glitch can be sufficiently deep for the function to touch y (last condition). Therefore, on line 6 hi' is set to mid, which satisfies the invariant. In fact, $\text{hi}' \in (\text{lo}, \text{hi})$ holds because $\text{lo} \prec \text{mid} \prec \text{hi}$, and $f(\text{mid}) \succ y$ implies that $f(\text{lo}) \preccurlyeq y \prec f(\text{hi}')$ holds. The last condition of the invariant is satisfied by this choice of hi', as proved at the beginning of this paragraph.

Otherwise, the control flow reaches the **else** body starting at line 8. At this point, we are not sure whether there is a glitch between mid and hi where the graph of the function reaches y or not. If there is only one glitch, its maximum width is small enough, and we know exactly where it starts, function `check_glitch` searches it float-by-float for a point z such that $f(z) \succcurlyeq y$. Whether it succeeded or not can be exerted from the value of b.

- $b = 0$: In this case, `check_glitch` could search the glitch, but it did not find any point in which the function was lower than or equal to y: setting $\text{hi}' = \text{mid}$ is correct. The first condition of the invariant is satisfied because of the post-condition of `split_point`, and the second condition holds because lo remained the same and, since we are into one of the **else** statements of the **if** at line 3, $f(\text{mid}) \succ y$. The third condition is also satisfied because the post-condition of `check_glitch` ensures that $\forall x \in [\text{mid}, \text{hi}] : f(x) \succ y$.
- $b = 1$: Function `check_glitch` succeeded in finding a value z such that $f(z) \preccurlyeq y$ inside the glitch. hi' is then set to z^+. $x_l \preccurlyeq \text{lo} \prec \text{hi}' \preccurlyeq x_u$ holds because the post-condition of `check_glitch` states that $\text{mid} \preccurlyeq z \preccurlyeq \text{hi}$, and both mid and hi are enclosed in x_l and x_u. The said post-condition also states that $\forall x \in (z, \text{hi}] : f(x) \succ y$: this maintains the last condition of the invariant, and implies that $f(z^+) \succ y$, which satisfies $f(\text{lo}) \preccurlyeq y \prec f(\text{hi}')$. This part of the post-condition of `check_glitch`, together with $\forall x \in [\text{hi}, x_u] : f(x) \succ y$, also implies that z is the rightmost point such that $f(z) \preccurlyeq y$. The **while** loop can therefore be broken, as the upper bound found until now cannot be refined further.
- $b = 2$: In this case the conditions that would allow `check_glitch` to search the glitch did not hold. Hence, lo and hi were left untouched, and the algorithm proceeds with the **else** block at line 14. Here, function `logsearch_ub` is called to find a suitable value for hi. Its post-condition allows us to distinguish between the two following cases.
  - $z \prec \text{hi}$: `logsearch_ub` found a value z such that $f(z) \succ^{d_M} y$: this assures us that there is no glitch in $[z, \text{hi}]$ deep enough for the graph of the function to reach y. Therefore, setting $\text{hi}' = z$ satis-

fies the last two terms of the invariant. Also, $\mathrm{hi}' \in (\mathrm{lo}, \mathrm{hi})$ holds because $z \in (\mathrm{mid}, \mathrm{hi})$.

$z = \mathrm{hi}$: In this case `logsearch_ub` could not find a safe value for hi, and we refrain from trying to refine the upper bound further. Since the loop is broken without changing anything, the invariant still holds.

Termination: In the previous paragraphs we proved that, either:

- the loop terminates with a **break** statement;
- the loop continues and one of lo' or hi' is set to a new value contained into the interval $(\mathrm{lo}, \mathrm{hi})$. In this case, the distance between lo and hi decreases at each iteration: the guard of the **while** loop, $\mathrm{hi} \succ^1 \mathrm{lo}$, will be eventually negated, terminating the loop.

Correctness: We have already proved that the invariant Inv holds after each iteration of the loop, and whenever it terminates with a **break** statement. Inv coincides with the "Correctness" part of the **Ensure** statement, which is therefore also proved.

If $\big(n_\mathrm{g} = 0 \ \vee \ \big(n_\mathrm{g} = 1 \wedge w_\mathrm{M} < t \ \wedge \ \big(\alpha = \alpha^f \vee \omega = \omega^f \vee \big(\alpha \prec^k \omega \wedge k \leq t\big)\big)\big)\big)$ holds, the post-condition of `check_glitch` assures us that either $b = 0$ or $b = 1$. Therefore, the only **break** statement that can be reached is the one at line 13. In this case, the same post-condition implies that $f(z) \preccurlyeq y$: setting $\mathrm{hi}' = z^+$ satisfies the "Precision" statement, because $f(\mathrm{hi}'^-) = f(z) \preccurlyeq y$. In all the other cases, the control flow can only reach lines 4, 6 or 10, which let the loop continue. If the loop is not terminated by the **break** statement at line 13, the same reasoning presented at the end of the proof of Algorithm 3 can be made: whenever the loop terminates because its guard is negated, we have $\#[\mathrm{lo}, \mathrm{hi}] = 2$. This, together with Inv, implies that $f(\mathrm{hi}^- = \mathrm{lo}) \preccurlyeq y$.

□

Theorem 4. *Function* `upper_bound`*, specified in* Algorithm 4*, satisfies its contract.*

Proof. Calls to functions `init` and `gallop_ub`, on lines 1 and 2 respectively, have the purpose of finding values for lo and hi suitable for the bisection phase. That is, as specified in the **Require** statement of function `bisect_ub` (Algorithm 7), $f(\mathrm{lo}) \preccurlyeq y \prec f(\mathrm{hi})$ and $\forall x \in [\mathrm{hi}, x_u] : f(x) \succ y$. Function `gallop_ub` is specular to its `lower_bound` counterpart: for more details about those functions, see the proof of Theorem 1.

The **if** statements before the call to `bisect_ub` on line 28 deal with the cases in which the said functions fail in their purpose. First, the value of $f(\mathrm{hi})$ is checked by the **if** statements on lines 4 and 7, leading to the following cases.

$f(\mathrm{hi}) \prec y$: In this case hi clearly does not satisfy the precondition of `bisect_ub` and, because of the post-condition of `gallop_ub` stated on line 3, $\mathrm{hi} = x_u$. Therefore, function `findhi_ub` is called on line 5 to further discern whether the equation $y = f(\mathbf{x})$ has no solution, or $\mathrm{hi} = x_u$ is inside a glitch and there is, in fact, a solution. Since function `findhi_ub` always sets u to a value satisfying p_6 (see Lemma 4), r can be set to 6, and function `upper_bound` can terminate.

$f(\mathrm{hi}) = y$: Again, according to the post-condition of `gallop_ub` we have $\mathrm{hi} = x_u$. Since hi is the highest value of interval $[x_l, x_u]$, setting $r = 9$ and $u = \mathrm{hi}$ is correct.

If control flow is not caught by the **if** statements described above, then $f(\mathrm{hi}) \succ y$ and, by the post-condition of `gallop_ub`, either of the following holds.

$x_u \succ \mathrm{hi}$: This implies that $f(\mathrm{hi}) \succ^{d_{\mathrm{M}}} y$. Therefore, no glitch in interval $[\mathrm{hi}, x_u]$ can be deep enough for function f to reach y: $\forall x \in [\mathrm{hi}, x_u] : f(x) \succ y$ holds.
$x_u = \mathrm{hi}$: In this case, $\forall x \in [\mathrm{hi}, x_u] \equiv [x_u, x_u] : f(x) \succ y$ trivially holds.

Therefore, hi satisfies the preconditions of `bisect_ub`.

Then, the value of f in lo is checked on line 10. If $f(\mathrm{lo}) \succ y$, lo is not suitable for bisection, and the body of the **if** statement on the said line tries to understand whether $y = f(\mathbf{x})$ has no solutions, or a solution is contained in a glitch. Note that the post-condition of `gallop_ub` entails that, in this case, $\mathrm{lo} = x_l$. The guard of the **if** statement on line 11 catches the two following cases.

$n_{\mathrm{g}} = 0$: In this case, function f is isotonic. Therefore, $\forall x \in [\mathrm{lo}, x_u] : f(x) \succcurlyeq f(\mathrm{lo})$. Together with $f(\mathrm{lo}) \succ y$, this implies that $\forall x \in [\mathrm{lo}, x_u] \equiv [x_l, x_u] : f(x) \succ y$. This lets us state that equation $y = f(\mathbf{x})$ has no solution, and setting $r = 5$ on line 12 is correct.
$f(\alpha) \succ^{d_{\mathrm{M}}} y$: This condition signifies that $f(\alpha)$ is too high for the graph of f to reach y inside a glitch after α. This, together with the definition of α and ω, brings the following conclusions:
- $\forall x \in [x_l, \alpha] : f(x) \succ y$ ($f(\mathrm{lo} = x_l) \succ y$ and isotonicity of f);
- $\forall x \in [\alpha, \omega] : f(x) \succ y$;
- $\forall x \in [\omega, x_u] : f(x) \succ y$ ($f(\omega) \succcurlyeq f(\alpha) \succ y$ and isotonicity of f).

Therefore, equation $y = f(\mathbf{x})$ has no solution in interval $[x_l, x_u]$, and we can set $r = 5$.

If none of the above conditions apply, there might be glitches deep enough to allow the function to take the value y somewhere in interval $[x_l, x_u]$. Function `check_glitch` is called with m = lo in line 14 to search such glitch, if it is only one and it is not too wide. The correctness of the said function is discussed in Lemma 5. The three values of b that this function can return are distinguished in the following **if** statements.

$b = 0$: By the post-condition of `check_glitch`, $\forall x \in [\mathrm{lo}, \mathrm{hi}] : f(x) \succ y$. As we previously noted, $\forall x \in [\mathrm{hi}, x_u] : f(x) \succ y$ holds at this point of the algorithm. This, together with the fact that $\mathrm{lo} = x_l$, implies that $\forall x \in [x_l, x_u] : f(x) \succ y$. Therefore, equation $y = f(\mathbf{x})$ has no solution, and we can set $r = 5$.
$b = 1$: The post-condition of `check_glitch` assures us that a value $z \in \mathbb{F}$ is returned, and $\mathrm{lo} \preccurlyeq z \preccurlyeq \mathrm{hi} \wedge \forall x \in (z, \mathrm{hi}] : f(x) \succ y \wedge f(z) \preccurlyeq y$ holds. We also know that $\forall x \in [\mathrm{hi}, x_u] : f(x) \succ y$ as stated before. This means that $\forall x \in (z, x_u] : f(x) \succ y$ and there is no solution after z. The algorithm must now distinguish whether z is a solution or not, and it does so starting from line 18.

$f(z) = y$: z is a solution to equation $y = f(\mathbf{x})$, and it is also the rightmost, for the reasons stated above. Therefore, setting $u = z$ and $r = 9$ is correct, and no further action is required, so the algorithm can **return**.

$f(z) \prec y$: We found the rightmost point where the graph of f crosses y without taking its value. Setting $u = z^+$ satisfies p_8, and the algorithm can terminate.

$b = 2$: In this last case, either $n_\mathrm{g} > 1$, or the glitch was too wide to be searched float-by-float. The only thing that can be done to refine the upper bound is setting u to the minimum between ω and hi. These values both satisfy p_7.
For hi, this is true because, as we noted before, at this point of the algorithm $\forall x \in [\mathrm{hi}, x_u] : f(x) \succ y$.
As for ω, because of the definition of α and ω, we are sure that f is isotonic outside interval $[\alpha, \omega]$ which, together with the fact that $f(\alpha) \preccurlyeq f(\omega)$ and $f(\mathrm{lo} = x_l) \succ y$, implies that $\forall x \in [\omega, x_u] : f(x) \succ y$.

If control flow does not enter the body of the **if** statement on line 10, then $f(\mathrm{lo}) \preccurlyeq y$, which satisfies the precondition of `bisect_ub`. Both lo and hi are suitable for the bisection phase, which is started on line 28 by calling `bisect_ub`. The correctness of this function is discussed in Lemma 6: after its call, the post-condition stated in the **Ensure** statement of Algorithm 7 is satisfied.

On line 29, a while loop approaches the solution of $y = f(\mathbf{x})$ by going backwards float-by-float. This loop is exactly specular to the one starting on line 20, and so is the proof of its correctness and termination. See Theorem 1 for more details, keeping in mind that the loop invariant is $\forall x \in [\mathrm{hi}, x_u] : f(x) \succ y$ in `upper_bound`.

The said invariant holds on line 33, where the value of $f(\mathrm{hi}^-)$ is checked, leading to the following cases.

$f(\mathrm{hi}^-) \prec y$: If $u = \mathrm{hi}$ is set, we have $f(u^-) \prec y \prec f(u)$. This, together with the **while**-loop invariant, satisfies p_8. u and r are set to the said values on line 34.

$f(\mathrm{hi}^-) = y$: The rightmost solution of $y = f(\mathbf{x})$ was found. u can be set to hi^- and r to 9, because p_9 is satisfied.

$f(\mathrm{hi}^-) \succ y$: This condition holds in the body of the **else** statement of line 37. The loop invariant, $\forall x \in [\mathrm{hi}, x_u] : f(x) \succ y$, assures that setting $u = \mathrm{hi}$ satisfies p_7. u and r are set accordingly on line 38.

The "Correctness" statement of this algorithm is therefore proved.

In order to prove the "Precision" statement, we assume that $f(x_l) \preccurlyeq y \preccurlyeq f(x_u)$ holds, together with either $n_\mathrm{g} = 0$ or $n_\mathrm{g} = 1 \wedge w_\mathrm{M} < t \wedge \big(\alpha = \alpha^f \vee \omega = \omega^f \vee (\alpha \prec^k \omega \wedge k \leq t)\big)$. This, together with the post-condition of `gallop_ub`, implies that after line 2 $f(\mathrm{lo}) \preccurlyeq y$ and $f(\mathrm{hi}) \succcurlyeq y$ hold, even if $\mathrm{lo} = x_l$ or $\mathrm{hi} = x_u$. Therefore, $f(\mathrm{lo})$ is always suitable for `bisect_ub`, and the **if** statement of line 10 cannot be entered. As for $f(\mathrm{hi})$, the following two cases must be distinguished.

$f(\mathrm{hi}) = y$: The body of the **if** statement on line 7 is entered, and $r = 9$ is returned, satisfying the "Precision" post-condition.

$f(\mathrm{hi}) \succ y$: $f(\mathrm{hi})$ satisfies the precondition of `bisect_ub`, that is called on line 28. After this call, according to the "Precision" post-condition of `bisect_ub`, $f(\mathrm{hi}^-) \preccurlyeq y$. The **while** loop on line 29 will therefore assign hi^- to hi, so that $f(\mathrm{hi}) \preccurlyeq y$. Only the bodies of the **if** statements on lines 33 and 35 can now be entered: r can only be set to 8 or 9, respectively. This satisfies the "Precision" post-condition of this algorithm.

□

Theorem 5. *In case $f\colon \mathbb{F} \to \mathbb{F}$ is isotonic, that is $n_\mathrm{g} = 0$, then, for any $[x_l, x_u] \in \mathcal{I}_\mathbb{F}$, d_M, w_M, α, ω, $f^\mathrm{i}\colon \mathbb{F} \to \mathbb{F}$, $s, t \in \mathbb{N}$, the call to the function* `upper_bound`$(f, y, [x_l, x_u], 0, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, f^\mathrm{i}, s, t)$ *in* Algorithm 4 *can be executed calling at most $2\log_2(\#[x_l, x_u]) + 4$ times the function f.*

Proof. The proof of this theorem is analogous to the proof of Theorem 2. It is therefore omitted.

Theorem 6. *In case the function $f\colon \mathbb{F} \to \mathbb{F}$ has small glitches, that is $n_\mathrm{g} > 0$ but $w_\mathrm{M} < t$ then, for any $[x_l, x_u] \in \mathcal{I}_\mathbb{F}$, d_M, α, ω, $f^\mathrm{i}\colon \mathbb{F} \to \mathbb{F}$, $s \in \mathbb{N}$, the call to function* `upper_bound`$(f, y, [x_l, x_u], n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, f^\mathrm{i}, s, t)$ *in* Algorithm 4 *can be executed calling at most $(w_\mathrm{M} + 1)\log_2(\#[x_l, x_u]) + 8$ times function f if the "Precision" clause holds, or at most $(s + 2)\log_2(\#[x_l, x_u]) - s + t + 7$ if it does not.*

Proof. The proof of this theorem is extremely similar to the one of Theorem 3: only the most significant differences will be discussed here. Moreover, as noted in the said proof, the control flow paths that present the highest number of calls to the function are those that include the call to `bisect_ub`. We will therefore restrict the discussion of this theorem to those paths.

If the "Precision" post-condition of Algorithm 4 holds then, on line 8 of function `bisect_ub`, function `check_glitch` is invoked. It performs the linear search of line 9, which behaves like the call to `findfmax` on line 22. Since in the worst case scenario this search is repeated in all iterations of the while loop of `bisect_ub`, the number of calls to f performed is $w_\mathrm{M}\log_2(\#[x_l, x_u])$, plus 4 extra calls to evaluate the guard of the **if** statement of line 1 of function `check_glitch`. Therefore, the overall number of calls to f of `upper_bound` is $(w_\mathrm{M} + 1)\log_2(\#[x_l, x_u]) + 8$ if the "Precision" condition holds.

Otherwise, function `logsearch_ub` is called for a maximum of s times on line 16 of algorithm `bisect_ub`. It calls f for at most $\log_2(\#[\mathrm{mid}, \mathrm{hi}])$ times. Since in every iteration of the main **while** loop of `bisect_ub` we have $\#[\mathrm{lo}, \mathrm{hi}] = 2\,\#[\mathrm{mid}, \mathrm{hi}]$, the actual number of function calls performed by `logsearch_ub` is $\log_2(\#[\mathrm{lo}, \mathrm{hi}]/2)$. Because `logsearch_ub` returns a value z which, in the worst case, is hi^-, we can consider as if it was always called on the whole initial $[\mathrm{lo}, \mathrm{hi}]$ interval. Since `logsearch_ub` is called for a maximum of s times, it leads to an additional number of calls to f at most $s(\log_2(\#[\mathrm{lo}, \mathrm{hi}]) - 1)$.

Then, after the call to `bisect_ub` has returned, the **while** loop of line 29 may be executed for up to t iterations, because the "Precision" condition does not hold.

The final count of the calls to function f sums up to

$$(s+2)\log_2(\#[x_l, x_u]) - s + t + 7.$$

□

B.3 Algorithms for Trigonometric Functions

Definition 8. *Let $x \in \mathbb{R}$, such that $\nexists n \in \mathbb{Z} \,.\, |x| = 2^n$ and $-f_{\max} < |x| < f_{\max}$ (i.e. x is not a power of 2, an infinity or subnormal). Let $[x] \in \mathbb{F}$ be a floating-point approximation of x with p significant digits and exponent e_x.*

Then, the rounding error for $[x]$ with respect to x when rounding towards zero, positive or negative infinity is $|\nabla(x)| \leq \text{ulp}(x) = 2^{e_x - p + 1}$. Moreover, with any of the roundToNearest rounding modes, the rounding error is $|\nabla(x)| \leq \frac{\text{ulp}(x)}{2} = 2^{e_x - p}$.

Lemma 1. *Let $k \in \mathbb{F}$ be a number in any of the IEEE 754 binary floating-point formats, with an exponent e_k. Let $x \in \mathbb{R}$, such that $\nexists n \in \mathbb{Z} \,.\, |x| = 2^n$ and $f_{\min} < |x| < f_{\max}$. Let $[x] \in \mathbb{F}$ be an approximation of x with p significant digits and exponent e_x, not necessarily in one of the said floating-point formats. All the operations are assumed to be correctly rounded.*

Then, the absolute error of the floating point multiplication between k and $[x]$ is $|\nabla(k \boxdot [x])| \leq 2^{e_k + e_x - p + 4}$.

Proof. The purpose of this lemma is to find a sufficiently tight upper bound for the rounding error of the said multiplication, keeping into account the initial rounding error of the second floating-point operand $[x]$. This can be written as the sum of the real number x plus its rounding error ($[x] = x + \nabla(x)$). Since we are assuming the multiplication is correctly rounded, it can also be written as the sum of its exact value plus its absolute error:

$$\begin{aligned} k \boxdot [x] &= [k(x + \nabla(x))] \\ &= [kx + k\,\nabla(x)] \\ &= kx + k\,\nabla(x) + \nabla(kx + k\,\nabla(x)) \\ &= kx + \nabla(k \boxdot [x]) \end{aligned}$$

We already have an upper bound for the rounding error of $[x]$, which is $|\nabla(x)| \leq 2^{e_x - p + 1}$. An upper bound for $|k|$ is given by $|k| < 2^{e_k + 1}$. Their product is thus bounded by the product of these two bounds.

An upper bound for $|kx + k\,\nabla(x)|$ can be found in the following way:

$$\begin{aligned} |kx + k\,\nabla(x)| &\leq kx + k|\,\nabla(x)| \\ &\leq 2^{e_k+1}2^{e_x+1} + 2^{e_k+1}2^{e_x-p+1} \\ &= 2^{e_k+e_x+2} + 2^{e_k+e_x-p+2} \\ &= 2^{e_k+e_x+2}(1 + 2^{-p}) \\ &< 2^{e_k+e_x+3} \end{aligned}$$

Thus, the exponent of $kx + k\,\nabla(x)$ is at most $2^{e_k+e_x+2}$, and

$$|\,\nabla(kx + k\,\nabla(x))| \leq \mathrm{ulp}(kx + k\,\nabla(x)) = 2^{e_k+e_x-p+3}.$$

We can now further derive the upper bound for the magnitude of the absolute error of the multiplication:

$$\begin{aligned}
|\,\nabla(k \boxdot [x])| &= |k\,\nabla(x) + \nabla(kx + k\,\nabla(x))| \\
&\leq k|\,\nabla(x)| + |\,\nabla(kx + k\,\nabla(x))| \\
&\leq 2^{e_k+e_x-p+2} + 2^{e_k+e_x-p+3} \\
&= 2^{e_k+e_x-p+2}(1+2) \\
&< 2^{e_k+e_x-p+4}
\end{aligned}$$

This proves that $|\,\nabla(k \boxdot [x])| \leq 2^{e_k+e_x-p+4}$.

Lemma 2. *If function f is quasi-monotonic over $[i.x_l, i.x_u]$, then the function* `branch_lb`*, specified in* Algorithm 9*, satisfies its contract.*

Proof. The **if** statement on line 1 prepares variables i_n_{g}, i_α and i_ω to be used as arguments for `lower_bound`. If there are no glitches in $[i.x_l, i.x_u]$, i_n_{g} is set to 0. Otherwise, the **else** statement ensures that $i.x_l \preccurlyeq i_\alpha \preccurlyeq i_\omega \preccurlyeq x_u$. Along with the preconditions of this algorithm, this satisfies all preconditions of `lower_bound`.

The only thing that remains to do, is to distinguish whether f is quasi-isotonic or quasi-antitonic in $[i.x_l, i.x_u]$. This is done by predicate quasi_isotonic $(i.k, p)$, by checking the value of $m = i.k \mod 4$: if $p = \mathrm{even_v}$, f is isotone when $m \in \{0, 3\}$; if $p = \mathrm{odd_v}$ f is isotone when $m \in \{1, 2\}$. If f is quasi-isotonic, `lower_bound` can be called normally. Otherwise, `lower_bound` is called on line 12 with $-f$ instead of f, $-y_u$ instead of y_l, and the inverse function f_i with the sign of its argument changed. Note that, if f is quasi-antitonic, then $-f$ is quasi-isotonic, which allows us to invoke `lower_bound` normally. Also, $-y_u$ must be passed as the lower bound for $\mathbf{y}$ because if $y_l \preccurlyeq y_u$, then $-y_u \preccurlyeq -y_l$.

Since $i.l$ and $i.r_l$ are set by `lower_bound`, they satisfy its post-condition. So $p^f_{i.r_l}(y_l, i.x_l, i.x_u, i.l)$ holds if f is isotonic in $[i.x_l, i.x_u]$, or $p^{-f}_{i.r_l}(-y_u, i.x_l, i.x_u, i.l)$ if f is antitonic in that interval. Note that, by $p^{-f}_{i.r_l}$, we intend predicate $p^f_{i.r_l}$ with all occurrences of f substituted with $-f$. Consequentially, also the **Ensure** statement of this algorithm holds. □

The algorithm listing and the proof of `branch_ub` are omitted, since they are very similar to those of `branch_lb`.

Theorem 7. *Function* `compute_bounds_trig`*, specified in* Algorithm 8*, satisfies its contract.*

Proof. The algorithm starts by identifying which branches of the function graph contain x_l and x_u. This operation is performed at lines 1 and 2, where function

Algorithm 8 Inverse propagation: `compute_bounds_trig`$(f, f^{\mathrm{i}}, p, y, [x_l, x_u], n_{\mathrm{gM}}, d_{\mathrm{M}}, w_{\mathrm{M}}, \alpha, \omega, g, s, t)$

Require: $f\colon \mathbb{F} \to \mathbb{F}$, $f^{\mathrm{i}}\colon \mathbb{F} \to \mathbb{F}$, $p \in \{\mathrm{even_v}, \mathrm{odd_v}, \mathrm{odd_c}\}$, $[x_l, x_u], [y_l, y_u] \in \mathcal{I}_{\mathbb{F}}$, $n_{\mathrm{gM}} \geq n^f_{\mathrm{gM}}$, $d_{\mathrm{M}} \geq d^f_{\mathrm{M}}$, $w_{\mathrm{M}} \geq w^f_{\mathrm{M}}$, $\alpha \preccurlyeq \alpha^f$, $\omega \succcurlyeq \omega^f$, $n_{\mathrm{gM}} > 0 \implies (x_l \preccurlyeq \alpha \preccurlyeq \omega \preccurlyeq x_u)$, $g, s, t \in \mathbb{N}$.

Ensure: $|I| \leq g$, $\bigcup_{i \in I} [i.x_l, i.x_u] = [x_l, x_u]$, $\forall i \in I : i.l, i.u \in \mathbb{F}$
$\wedge \left(i.r_l \in \{0,1,2,3,4\} \implies p^f_{i.r_l}(y_l, i.x_l, i.x_u, i.l) \vee p^{-f}_{i.r_l}(-y_u, i.x_l, i.x_u, i.l) \right)$
$\wedge \left(i.r_u \in \{5,6,7,8,9\} \implies p^f_{i.r_u}(y_u, i.x_l, i.x_u, i.u) \vee p^{-f}_{i.r_u}(-y_l, i.x_l, i.x_u, i.u) \right)$.

$k_l := \mathtt{geq_tonicity_change}(p, \lceil x_l \frac{2}{\pi} \rceil)$ ▷ $k_l = \min\{k \in \mathbb{Z} \mid k \geq \lceil x_l \frac{2}{\pi} \rceil \wedge \mathrm{tonicity_change}(k, p)\}$
$k_u := \mathtt{geq_tonicity_change}(p, \lceil x_u \frac{2}{\pi} \rceil)$ ▷ $k_u = \min\{k \in \mathbb{Z} \mid k \geq \lceil x_u \frac{2}{\pi} \rceil \wedge \mathrm{tonicity_change}(k, p)\}$
if $\mathrm{even}(p)$ **then** $k_{\mathrm{c}} := 0$
else $k_{\mathrm{c}} := -1$
end if
if $g = 1$ **then** $g_{\mathrm{l}} := 1$; $g_{\mathrm{r}} := 0$; $k_{\mathrm{c}} := k_u$
else if $x_u \preccurlyeq \left[k_{\mathrm{c}} \frac{\pi}{2}\right]_{\downarrow}$ **then** $g_{\mathrm{l}} = g$; $g_{\mathrm{r}} = 0$
else if $x_l \succcurlyeq \left[k_{\mathrm{c}} \frac{\pi}{2}\right]_{\uparrow}$ **then** $g_{\mathrm{l}} := 0$; $g_{\mathrm{r}} := g$
else $g_{\mathrm{l}} := \left\lfloor \frac{g}{2} \right\rfloor$; $g_{\mathrm{r}} := g - g_{\mathrm{l}}$
end if
if $g_{\mathrm{l}} > 0$ **then**
    $i.x_l := x_l$; $i.x_u := \min\left\{x_u, \left[k_l\right]_{\downarrow}\right\}$
    $(i.l, i.r_l) := \mathtt{branch_lb}(f, f^{\mathrm{i}}, p, [y_l, y_u], [i.x_l, i.x_u], k_l, n_{\mathrm{gM}}, d_{\mathrm{M}}, w_{\mathrm{M}}, \alpha, \omega, s, t)$
    $k_{\mathrm{l}_u} := \min\{k_u, k_{\mathrm{c}}\}$
    $i.k := \max\{k_l, k_{\mathrm{l}_u} - 2(g_{\mathrm{l}} - 1)\}$
    if $i.r_l = 0 \wedge k_l < i.k$ **then**
        $i.l := i.x_u$; $i.r_l := 2$
    end if
    $c_{x_l} := \max\left\{x_l, \left[(i.k - 2)\frac{\pi}{2}\right]_{\uparrow}\right\}$
    $i.x_u := \min\left\{x_u, \left[i.k \frac{\pi}{2}\right]_{\downarrow}\right\}$
    $(i.u, i.r_u) := \mathtt{branch_ub}(f, f^{\mathrm{i}}, p, [y_l, y_u], [c_{x_l}, i.x_u], i.k, n_{\mathrm{gM}}, d_{\mathrm{M}}, w_{\mathrm{M}}, \alpha, \omega, s, t)$
    if $i.r_u = 5 \wedge k_l < i.k$ **then**
        $i.u := i.x_l$; $i.r_u := 7$
    end if
    $I.\mathtt{add}(i)$
    $i.k := i.k + 2$
    while $i.k < k_{\mathrm{l}_u}$ **do**
        $i.x_l := {i.x_u}^+$; $i.x_u := \min\left\{\left[i.k \frac{\pi}{2}\right]_{\downarrow}, x_u\right\}$
        $(i.l, i.r_l) := \mathtt{branch_lb}(f, f^{\mathrm{i}}, p, [y_l, y_u], [i.x_l, i.x_u], i.k, n_{\mathrm{gM}}, d_{\mathrm{M}}, w_{\mathrm{M}}, \alpha, \omega, s, t)$
        $(i.u, i.r_u) := \mathtt{branch_ub}(f, f^{\mathrm{i}}, p, [y_l, y_u], [i.x_l, i.x_u], i.k, n_{\mathrm{gM}}, d_{\mathrm{M}}, w_{\mathrm{M}}, \alpha, \omega, s, t)$
        $I.\mathtt{add}(i)$
        $i.k := i.k + 2$
    end while
end if

if $g_\mathrm{r} > 0$ **then**
  $i.k := \max\{k_l, k_\mathrm{c} + 2\}$
  $i.x_u := \max\left\{{x_l}^-, \left[(i.k-2)\frac{\pi}{2}\right]_\downarrow\right\}$
  while $i.k < k_u \wedge g_\mathrm{r} > 1$ **do**
    $i.x_l := {i.x_u}^+$; $i.x_u := \min\left\{x_u, \left[i.k\frac{\pi}{2}\right]_\downarrow\right\}$
    $(i.l, i.r_l) := \texttt{branch_lb}(f, f^\mathrm{i}, p, [y_l, y_u], [i.x_l, i.x_u], i.k, n_\mathrm{gM}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, s, t)$
    $(i.u, i.r_u) := \texttt{branch_ub}(f, f^\mathrm{i}, p, [y_l, y_u], [i.x_l, i.x_u], i.k, n_\mathrm{gM}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, s, t)$
    $I.\texttt{add}(i)$
    $i.k := i.k + 2$
    $g_\mathrm{r} := g_\mathrm{r} - 1$
  end while
  $i.x_l := {i.x_u}^+$; $i.x_u := \min\left\{x_u, [i.k\frac{\pi}{2}]_\downarrow\right\}$
  $(i.l, i.r_l) := \texttt{branch_lb}(f, f^\mathrm{i}, p, [y_l, y_u], [i.x_l, i.x_u], i.k, n_\mathrm{gM}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, s, t)$
  if $i.r_l = 0 \wedge i.k < k_u$ **then**
    $i.l := i.x_u$; $i.r_l := 2$
  end if
  $c_{x_l} := \max\left\{x_l, \left[(k_u-2)\frac{\pi}{2}\right]_\uparrow\right\}$; $i.x_u := x_u$
  $(i.u, i.r_u) := \texttt{branch_ub}(f, f^\mathrm{i}, p, [y_l, y_u], [c_{x_l}, i.x_u], k_u, n_\mathrm{gM}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, s, t)$
  if $i.r_u = 5 \wedge i.k < k_u$ **then**
    $i.u := c_{x_l}$; $i.r_u := 7$
  end if
  $I.\texttt{add}(i)$
end if

Algorithm 9 Inverse propagation: $\texttt{branch_lb}(f, f^\mathrm{i}, p, y, [i.x_l, i.x_u], i.k, n_\mathrm{gM}, d_\mathrm{M}, w_\mathrm{M}, \alpha, \omega, s, t)$

Require: $f\colon \mathbb{F} \to \mathbb{F}$, $f^\mathrm{i}\colon \mathbb{F} \to \mathbb{F}$, $p \in \{\mathrm{even_v}, \mathrm{odd_v}, \mathrm{odd_c}\}$, $y \in \mathbb{F}$, $[i.x_l, i.x_u], [y_l, y_u] \in \mathcal{I}_\mathbb{F}$, $i.k \in \mathbb{Z}$, $n_\mathrm{gM} \geq n^f_\mathrm{gM}$, $d_\mathrm{M} \geq d^f_\mathrm{M}$, $w_\mathrm{M} \geq w^f_\mathrm{M}$, $\alpha \preccurlyeq \alpha^f$, $\omega \succcurlyeq \omega^f$, $n_\mathrm{gM} > 0 \implies (\alpha \preccurlyeq \omega)$, $s, t \in \mathbb{N}$.

Ensure: Correctness: $i.l \in \mathbb{F}$,

$$i.r_l \in \{0,1,2,3,4\} \implies \left(p^f_{i.r_l}(y_l, i.x_l, i.x_u, i.l) \vee p^{-f}_{i.r_l}(-y_u, i.x_l, i.x_u, i.l)\right).$$

if $n_\mathrm{gM} = 0 \vee i.x_l \succ \omega \vee i.x_u \prec \alpha$ **then**
  $i_n_\mathrm{g} := 0$
else
  $i_n_\mathrm{g} := n_\mathrm{gM}$
  $i_\alpha := \max\{\alpha, i.x_l\}$
  $i_\omega := \min\{\omega, i.x_u\}$
end if
if quasi_isotonic$(i.k, p)$ **then**
  $(i.l, i.r_l) := \texttt{lower_bound}(f, y_l, [i.x_l, i.x_u], i_n_\mathrm{g}, d_\mathrm{M}, w_\mathrm{M}, i_\alpha, i_\omega, f^\mathrm{i}, s, t)$
else
  $(i.l, i.r_l) := \texttt{lower_bound}(-f, -y_u, [i.x_l, i.x_u], i_n_\mathrm{g}$
    $d_\mathrm{M}, w_\mathrm{M}, i_\alpha, i_\omega, f^\mathrm{i} \circ (-\mathrm{id}), s, t)$
end if

`geq_tonicity_change` is called. The lowest integer that, multiplied by $\frac{\pi}{2}$, gives a number greater than x_l where the monotonicity of the function changes is assigned to k_l. The same is done for x_u, and the value is assigned to k_u.

On line 3 a value for k_c is chosen. The approximations of $k_\mathrm{c}\frac{\pi}{2}$ will be used to split $[x_l, x_u]$ in two intervals $\left[x_l, \left[k_\mathrm{c}\frac{\pi}{2}\right]_\downarrow\right]$ and $\left[\left[k_\mathrm{c}\frac{\pi}{2}\right]_\uparrow, x_u\right]$, which will be referred to as "left" and "right" intervals respectively. Branches of the function nearest to $k_\mathrm{c}\frac{\pi}{2}$ will not be included into the two large intervals near the bounds of the domains, and will be processed individually. If f changes monotonicity in even multiples of $\frac{\pi}{2}$ (such as the cosine), $k_\mathrm{c} = 0$. Otherwise, k_c is set to -1: in this way the branch between $-1\frac{\pi}{2}$ and $\frac{\pi}{2}$ will be the leftmost one in the right interval.

Now, values for g_l and g_r are chosen. They are, respectively, the number of sub-intervals in which the left and right intervals will be divided. The **if** statement at line 6 distinguishes among the four cases listed below.

$g = 1$: Only one interval must be returned, corresponding to $[x_l, x_u]$, but refined in the outward branches of the function. In this case, g_l is set to 1, and g_r to 0. In this way, the body of the **if** statement on line 11 will take care of the said interval, as it would normally do with the leftmost interval. Also, $k_c := k_u$ is set, since the above mentioned code uses k_c to chose the upper bound of the said interval.

$g > 1 \wedge x_u \preccurlyeq \left[k_\mathrm{c}\frac{\pi}{2}\right]_\downarrow$: $[x_l, x_u]$ lies entirely into the left interval. Therefore, g_l must be set to g and g_r to 0.

$g > 1 \wedge x_l \succcurlyeq \left[k_\mathrm{c}\frac{\pi}{2}\right]_\uparrow$: $[x_l, x_u]$ lies entirely into the right interval. Therefore, g_l must be set to 0 and g_r to g.

$g > 1 \wedge x_l \prec \left[k_\mathrm{c}\frac{\pi}{2}\right]_\uparrow \wedge x_u \succ \left[k_\mathrm{c}\frac{\pi}{2}\right]_\downarrow$: $[x_l, x_u]$ must be divided into left and right intervals. g_l is set to $\left\lfloor\frac{g}{2}\right\rfloor$ and g_r is set to $g - g_\mathrm{l}$. This way, half of the sub-interval will be in the left interval, and half in the right one. Note that, if the function has less than g_l monotonic branches in the left interval or less then g_r in the right one, the interval will be only split in as many intervals as the branches.

In all four cases, $g_\mathrm{l} + g_\mathrm{r} = g$.

The purpose of the body of the **if** statement on line 11 is to divide the left interval $\left[x_l, \min\left\{x_u, \left[k_\mathrm{c}\frac{\pi}{2}\right]_\downarrow\right\}\right]$ into the appropriate sub-intervals, and refine them by calling `branch_lb` and `branch_ub`. First, the leftmost sub-interval is generated. Each sub-interval is identified by its initial bounds $i.x_l$ and $i.x_u$, by the refined bounds $i.l$ and $i.u$, by $i.r_l$ and $i.r_u$ (the predicates that hold after calling `lower_bound` and `upper_bound`), and by $i.k$, which identifies the branch (or one of the branches) of the function in which the sub-interval is contained. That is, it defines the interval $\left[\left[(i.k-2)\frac{\pi}{2}\right]_\uparrow, \left[i.k\frac{\pi}{2}\right]_\downarrow\right]$.

If the number of monotonic branches in the left interval is lower than or equal to g_l, then this sub-interval will correspond to a single branch; it will include multiple branches otherwise. On line 12 $i.x_l$ and $i.x_u$ are set to the bounds of the leftmost branch, $\left[x_l, \min\left\{x_u, \left[k_l\right]_\downarrow\right\}\right]$. Because of how k_l was set, this interval

is part of the leftmost monotonic branch of the function, and it covers the end of it if $x_u \geq \left[k_l\right]_{\downarrow}$. Then since f is monotonic in $[i.x_l, i.x_u]$, `branch_lb` can be called and $i.l$ and $i.r_l$ are set, satisfying $p^f_{i.r_l}(y_l, i.x_l, i.x_u, i.l)$ if f is isotonic in that interval, or $p^{-f}_{i.r_l}(-y_u, i.x_l, i.x_u, i.l)$ if it is antitonic, according to the post-condition of `branch_lb`.

The left interval will end with the branch defined by k_c, or with x_u if $x_u \prec \left[k_\mathrm{c}\frac{\pi}{2}\right]_{\downarrow}$. On line 14, k_{l_u} is set accordingly to $\min\{k_u, k_\mathrm{c}\}$. Then, on line 15 the rightmost branch of the left interval is chosen: it will end in $\left[i.k\frac{\pi}{2}\right]_{\downarrow}$. Since the rightmost $g_\mathrm{l} - 1$ branches will be treated individually by the **while** loop on line 27, $i.k$ is set to $\max\{k_l, k_{\mathrm{l}_u} - 2(g_\mathrm{l} - 1)\}$. $g_\mathrm{l} - 1$ is multiplied by 2 because each branch has a length of π (two times $\frac{\pi}{2}$), and subtracted from k_{l_u}, which identifies the rightmost possible branch in the left interval.

Before passing to the refinement of the upper bound of the sub-interval, the **if** statement on line 16 checks whether $i.r_l$, the predicate number returned by `lower_bound`, is 0. If $k_l = i.k$, it means this sub-interval contains a single branch, and no more work is needed, because we know there are no solutions to $y_l = f(\mathbf{x})$ in this sub-interval. If also $k_l < i.k$ holds, there are more branches after $i.x_u$ in this sub-interval. According to p_0 we have $\forall x \in [i.x_l, i.x_u] : y \succ f(x)$ but we know nothing about the other branches. Since $i.x_u$ will be later set to the upper bound of this sub-interval, $i.l$ must be set to the current value of $i.x_u$, and $i.r_l$ to 2. This way, we have $\forall x \in [i.x_l, i.l] : y \succ f(x)$, and either p^f_2 or p^{-f}_2 is satisfied even if the value of $i.x_u$ is later increased.

Then, the rightmost branch of the leftmost sub-interval is refined. On line 19, c_{x_l} is set to the maximum between x_l (remain inside the domain, in case this sub-interval contains only one branch), and $\left[(i.k-2)\frac{\pi}{2}\right]_{\uparrow}$. This is the lower bound of the rightmost branch of this sub-interval (note that $i.k_u$ was previously set to the corresponding value). Similarly, $i.x_u$ is set to the minimum between x_u and $\left[i.k\frac{\pi}{2}\right]_{\downarrow}$ on line 20. Now, f is monotonic over $[c_{x_l}, i.x_u]$, and we are ready to refine the upper bound of this sub-interval, by calling `branch_ub`.

The **if** statement of line 22 checks whether $i.r_u = 5$. Again, if $k_l = i.k$ nothing needs to be done. If $k_l < i.k$, we cannot return $i.r_u = 5$, because we do not know if there are solutions or not in the branches we did not analyze. Therefore, $i.u$ is set to $i.x_l$ and $i.r_u$ to 7. Since the post-condition of `branch_ub` ensures that p_5 holds, we have $\forall x \in [i.x_l, i.x_u] : y \prec f(x)$: then also $\forall x \in [i.u, i.x_u] : y \prec f(x)$ holds, which satisfies p^f_7 or p^{-f}_7.

At this point, both bounds of the leftmost interval have been refined, and i, with the current values of its fields, can be added to I. Then we can pass to the next sub-interval, identified by the current $i.k + 2$. $i.k$ is updated accordingly on line 26.

The **while** loop starting on line 27 refines all other sub-intervals in the left interval, which all correspond to a single monotonic branch of the function. We

will now prove that loop invariant

$$Inv \equiv \bigcup_{h \in I} [h.x_l, h.x_u] = [x_l, i.x_u] \tag{13}$$

$$\wedge\, \forall h \in I : h.l, h.u \in \mathbb{F} \tag{14}$$

$$\wedge\, h.r_l \in \{0,1,2,3,4\} \implies \begin{array}{c} p^f_{h.r_l}(y_l, h.x_l, h.x_u, h.l) \\ \vee p^{-f}_{h.r_l}(-y_u, h.x_l, h.x_u, h.l) \end{array} \tag{15}$$

$$\wedge\, h.r_u \in \{5,6,7,8,9\} \implies \begin{array}{c} p^f_{h.r_u}(y_u, h.x_l, h.x_u, h.u) \\ \vee p^{-f}_{h.r_u}(-y_l, h.x_l, h.x_u, h.u) \end{array} \tag{16}$$

$$\wedge\, i.x_u = \min\left\{ \left[(i.k-2)\frac{\pi}{2}\right]_{\downarrow}, x_u \right\} \tag{17}$$

holds before and after each iteration of the loop, unless the loop guard is false before the loop is entered the first time. In this latter case, clause (17) does not hold, but the other do.

Before starting with the proof of the loop, we need to prove the following fact: right after lines 20 and 28,

$$x_u \preccurlyeq \left[i.k\frac{\pi}{2}\right]_{\downarrow} \implies i.k + 2 \geq k_{l_u}. \tag{18}$$

Because of how k_u was set on line 2, we also have $\left[(k_u-2)\frac{\pi}{2}\right]_{\uparrow} \preccurlyeq x_u$. This implies $\left[(k_u-2)\frac{\pi}{2}\right]_{\uparrow} \preccurlyeq \left[i.k\frac{\pi}{2}\right]_{\downarrow}$ and $k_u - 2 \leq i.k$. Line 14 ensures that $k_{l_u} \leq k_u$: we have $k_{l_u} - 2 \leq k_u - 2 \leq i.k$, or $i.k + 2 \geq k_{l_u}$.

Initialization: Before of the first iteration of the loop, I contains only one interval i. Since $i.x_l$ was set to x_l on line 12 and left untouched, and also $i.x_u$ was not changed after the interval was added to I, clause (13) of the invariant is satisfied. As for $i.l$, it satisfies predicate $p_{i.r_l}$ because, either:

- $i.r_l \neq 0$: it was set by `branch_lb` and left untouched. Even if $i.x_u$ has been possibly increased after `branch_lb` was called, all predicates $p_{i.r_l}$ with $i.r_l \neq 0$ still hold, because they do not imply anything so specific about interval $(i.l, i.x_u]$ to be invalidated by increasing $i.x_u$.
- $i.r_l = 0$: the **if** statement on line 16 changed $i.l$ and $i.r_l$ so that $p_{i.r_l}$ still holds, as was proved before.

A very similar reasoning can be done for $i.u$ and $i.r_u$, distinguishing whether `branch_ub` returned $i.r_u = 5$ or not. Therefore, clauses (15) and (16) of the invariant are also proved.

On line 20, $i.x_u$ was set to $i.x_u := \min\left\{x_u, \left[i.k\frac{\pi}{2}\right]_{\downarrow}\right\}$ and, since $i.k$ was incremented by 2 on line 26, clause (17) is satisfied.

Maintenance: If the loop body was entered, the guard condition is true: so $i.k < k_{l_u}$. This, together with implication (18), assures us that $x_u \succ \left[i.k\frac{\pi}{2}\right]_{\downarrow}$. Because of clause (17) of the invariant, $x_u = \left[(i.k-2)\frac{\pi}{2}\right]_{\downarrow}$, and ${x_u}^+ =$

$\left[(i.k-2)\frac{\pi}{2}\right]_{\uparrow}$. $i.x_l$ is set to this value on line 28. On the same line, $i.x_u$ is set to the minimum between x_u and $\left[i.k\frac{\pi}{2}\right]_{\downarrow}$.

Therefore, $[i.x_l, i.x_u]$ is contained in a single monotonic branch, and functions `branch_lb` and `branch_ub` can be called to refine the bounds of this sub-interval. The post-conditions of the said functions satisfy clauses (15) and (16) of the invariant. Sub-interval i is then added to I. Since on line 28 $i.x_l$ was set to the successor of the previous value of $i.x_u$, and because the invariant ensures that previously $\bigcup_{h\in I}[h.x_l, h.x_u] = [x_l, i.x_u]$, adding the new interval to I leaves clause (13) satisfied. Eventually, after line 32 clause (17) also holds.

Termination: On line 15, i_k was set to a value such that $k_{\mathrm{l}_u} - i.k \leq 2(g_{\mathrm{l}} - 1)$. Since $i.k$ is incremented by 2 at the end of the loop body, it reaches k_{l_u} after at most $g_{\mathrm{l}} - 1$ iterations: the guard then evaluates to false, and the loop terminates.

Since the loop iterates at most $g_{\mathrm{l}} - 1$ times, the number of sub-intervals it adds to I is lower than or equal to $g_{\mathrm{l}} - 1$. Together with the fact that the leftmost sub-interval is always added to I before the loop, this implies $|I| \leq g_{\mathrm{l}}$.

At this point, if $k_u \leq k_{\mathrm{c}}$, then $k_{\mathrm{l}_u} = k_u$ and the last sub-interval ended in x_u: $i.x_u = x_u$ and clause $\bigcup_{i\in I}[i.x_l, i.x_u] = [x_l, x_u]$ of the post-condition is proved. In this case, $x_u \preccurlyeq \left[k_{\mathrm{c}}\frac{\pi}{2}\right]_{\downarrow}$ and $g_{\mathrm{r}} = 0$ according to line 7: the **if** statement of line 35 is not entered, and the algorithm terminates. Since there are no more sub-intervals to be analyzed, $|I| \leq g_{\mathrm{l}} = g$ and the loop invariant satisfy the rest of the post-condition.

Otherwise, if x_u has not been reached yet, the **if** body of line 35 processes the remaining sub-intervals. First, it splits the first $g_{\mathrm{r}} - 1$ branches of f and refines them with `branch_lb` and `branch_ub`. The loop on line 38 terminates when either the rightmost branch is refined, or $g_{\mathrm{r}} = 1$. Then, a sub-interval made of the remaining branches is processed in a way similar to the one of the leftmost interval. Since this part of the algorithm is similar to the **if** statement of line 11, the proof of its correctness is omitted.

After the body of this last **if** statement is executed, a number of sub-intervals less or equal than g_{r} has been added to I. Together with the fact that $|I| \leq g_{\mathrm{l}}$ held before, this implies that at the end $|I| \leq g$. The fact that it processes all branches between $\left[k_{\mathrm{c}}\frac{\pi}{2}\right]_{\uparrow}$ and x_u by calling `branch_lb` and `branch_ub`, and it corrects their output for the rightmost interval if needed, proves the rest of the post-condition of this algorithm. □